Technical Report 1162

Exploring the Interaction of Implicit and Explicit Processes to Facilitate Individual Skill Learning

Ron Sun
Rensselaer Polytechnic Institute

Robert C. Mathews
Louisiana State University

May 2005

20050713 103

United States Army Research Institute
for the Behavioral and Social Sciences

U.S. Army Research Institute
for the Behavioral and Social Sciences

A Directorate of the Department of the Army
Deputy Chief of Staff, G1

ZITA M. SIMUTIS
Director

NOTICES

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YYYY)* May 2005	2. REPORT TYPE Final		3. DATES COVERED *(From - To)* September 2000 – February 2005
4. TITLE AND SUBTITLE Exploring the Interaction of Implicit and Explicit Processes to Facilitate Individual Skill Learning		5a. CONTRACT NUMBER DASW01-00-K-0012	
		5b. PROGRAM ELEMENT NUMBER 611102	
6. AUTHOR(S) Ron Sun (Rensselaer Polytechnic Institute) and Robert C. Mathews (Louisiana State University)		5c. PROJECT NUMBER B74F	
		5d. TASK NUMBER 2902	
		5e. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Rensselaer Polytechnic Institute 110 8th Street Troy, NY 12180-3590		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Institute for the Behavioral and Social Sciences 2511 Jefferson Davis Highway Arlington, VA 22202-3926		10. MONITOR'S ACRONYM(S) ARI 11. MONITOR'S REPORT NUMBER(S) Technical Report 1162	

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
COR: Dr. Paul Gade

14. ABSTRACT
This work advances basic research in the areas of learning and training. One product is a conceptual framework, which addresses the ways explicit and implicit knowledge interact to produce skills. This framework suggests that human performance may be controlled by either a subconceptual knowledge base (the implicit mode) or application of a symbolic conceptual model (the explicit mode). A computational cognitive architecture, CLARION, significantly different from other existing cognitive architectures, is developed in this work to capture a range of data related to the interaction. It helps us to explain (and eventually to predict) training and learning processes. The results of the experiments support the theory of the interactions of implicit and explicit learning processes during skill acquisition. The outcomes (data, models, and theories) provide a more detailed, clearer and more comprehensive perspective on skill learning.

15. SUBJECT TERMS
skill acquisition, implicit learning, cognition, computational modeling, cognitive modeling, process control,

SECURITY CLASSIFICATION OF:			19. LIMITATION OF ABSTRACT	20. NUMBER OF PAGES	21. RESPONSIBLE PERSON Ellen Kinzer Technical Publication Specialist 703-602-8047
16. REPORT Unclassified	17. ABSTRACT Unclassified	18. THIS PAGE Unclassified	Unlimited	85	

Technical Report 1162

Exploring the Interaction of Implicit and Explicit Processes to Facilitate Individual Skill Learning

Ron Sun
Rensselaer Polytechnic Institute

Robert C. Mathews
Louisiana State University

Research and Advanced Concepts Office
Paul A. Gade, Chief

FOREWORD

Any high level of skill depends on both conceptual (explicit) and subconceptual (experiential or implicit) knowledge. However, experts are often only aware of their explicit conceptual knowledge. Experientially acquired implicit knowledge is more akin to pattern recognition. For example, when you recognize a friend's face, you instantly know who the person is, but you may not be aware of what cues or features are being used to recognize him/her. This lack of awareness of essential implicit experiential knowledge creates serious challenges for training and learning in the military as well as civilian context. The purpose of this document is to describe the work on one U.S. Army Research Institute for the Behavioral and Social Sciences (ARI) research project that seeks to understand how implicitly acquired knowledge from experience interacts with explicitly acquired knowledge (mental models) leading to better computational theory/models of human learning. This research incorporated both experimental and theoretical work, culminating in the refinement of the CLARION computational model of implicit and explicit learning.

MICHELLE SAMS
Technical Director

ACKNOWLEDGMENTS

EXPLORING THE INTERACTION OF IMPLICIT AND EXPLICIT PROCESSES TO FACILITATE INDIVIDUAL SKILL LEARNING

EXECUTIVE SUMMARY

Research Requirement:

Improving the speed and quality of training in complex skills continues to be a major need in the military. A basic problem for the Army is how to ensure that novices in a Military Occupational Specialty (MOS) move quickly to more advanced performance (and perhaps to expertise) as a result of their training. In addition, most training focuses on teaching conceptual (explicit) knowledge rather than setting up the opportunity for substantial experiential (implicit) knowledge. While this may be appropriate for some specialties, some other specialties involve working with complex systems that are better learned initially through extensive experience (implicit learning) than with lectures or textbook lessons (explicit learning). In many situations it is not clear what is the optimal mix of implicit (hands-on) training methods and explicit instructions. The particular mix of training not only affects acquired level of expertise, but also the relative speed and accuracy of decisions involved in performing a complex skill.

Procedure:

Two series of experiments were conducted in two different task domains: process control and artificial grammar learning. Both tasks involve learning a system that operates according to complex, difficult-to-learn rules. Laboratory experiments using college students as subjects were conducted.

The process control task involved learning to control the temperature of a simulated nuclear reactor by controlling the number of fuel pellets. The appropriate number of pellets to use depended on the current temperature of the reactor, sometimes creating counterintuitive situations where increasing the number of fuel pellets decreases temperature. Also, a noise element was included in the formula making the results somewhat uncertain over trials. This task is known to be difficult to learn and difficult to explain how one accomplishes the task when it is learned.

In the artificial grammar experiments, participants learned to generate poison can labels on a computer simulation of a starship that had been invaded by enemy agents. Identification of poison food labels requires learning to identify sequences of letters generated by a finite state grammar. Training for this task in different experiments included memorizing a diagram of the grammar (explicit training), memorizing cases (implicit training), or an integrated (implicit and explicit) training that involves tracing cases through the grammar diagram.

A computational cognitive architecture, CLARION, markedly different from other existing cognitive architectures, is developed in this work to capture a range of quantitative data related to the interaction of implicit and explicit learning. We carry out simulation experiments

in the domains of process control tasks, artificial grammar learning tasks, as well as many other tasks, further explicating the interaction between implicit and explicit processes.

Findings:

- With only one type of training we found slow but accurate responding when explicitly trained, and fast but less accurate responding when implicitly trained.
- Integrating or mixing types of training generally produced more accurate performance than implicit training and faster performance than strict explicit training.
- When exposed to both types of training, participants showed a tendency to prefer using the implicit mode to perform the task.
- In the case of strict explicit followed by implicit training (a pattern that is common in many training situations, e.g., explicit schooling followed by field training) we noticed a loss of accuracy as learners shifted toward the implicit mode after training.
- We were able to obtain the best of both worlds—fast and accurate responding—by incorporating an animated version of the task constraints (a diagram of the grammar) that indicated how current exemplars fit into the model during practice.
- In the process control domain, reflection during task performance interferes with learning.
- However, reflection following short periods of practice can be beneficial early in learning.
- Implicitly acquired knowledge can be much more flexible than existing research suspected.
- There are large individual differences in the ability to learn the process control task. Potentially, we can facilitate learning in these "failing" participants by instructing them on what to focus on when they reflect.
- A simple hint consisting of providing three correct responses to particular task situations greatly increased learning.
- We think training can be accelerated and this post-training drop eliminated by using the explicit conceptual knowledge to prime rather than compete with implicit learning.

Utilization of Findings:

These findings can be used to develop training principles that enhance training and lead to fast accurate decisions.

EXPLORING THE INTERACTION OF IMPLICIT AND EXPLICIT PROCESSES TO FACILITATE INDIVIDUAL SKILL LEARNING

CONTENTS
_____

CONTENTS (continued)

Page

LIST OF FIGURES

LIST OF TABLES

Exploring the Interaction of Implicit and Explicit Processes to Facilitate Individual Skill Learning

In the skill acquisition literature, the role of implicit learning in skill acquisition and the distinction between implicit and explicit learning have been recognized. However, although implicit learning has been actively investigated, the interaction between the implicit and the explicit has not been sufficiently explored. Research has been focused on showing the lack of explicit learning in various learning settings. Similar oversight is also evident in most computational simulation models of implicit learning. Despite the lack of studies of interaction, there is mounting evidence that it is difficult to find a situation in which only one type of learning is engaged. Our review of existing data (Sun 2002) indicated that in most situations, both types of learning are involved, with varying amounts of contributions from each. Therefore, in this research, we focus on studying the interaction between implicit and explicit processes in skill acquisition and how this interaction may be used to enhance training.

In this report, we first present a simulation examination of process control data, which led to some interesting initial hypotheses concerning implicit vs. explicit processes, which necessitated validation with human experiments. Specifically, the simulation explicates the interaction between the implicit and explicit learning processes in skill acquisition and highlights the interaction between the two types of processes and its various effects on learning (including the synergy effect). This simulation utilizes an integrated model (named CLARION) of skill learning that takes into account both implicit and explicit processes; moreover, it embodies a bottom-up approach (first learning implicit knowledge and then explicit knowledge on its basis) towards skill learning. The simulation shows that this approach accounts for various effects in the process control task data that have previously reported in the literature. Now the question is how we verify the chief hypothesis of this simulation: The interaction between implicit and explicit knowledge is the key to understanding and enhancing human skill learning. In the sections that follow this simulation section, we address this question, chiefly through human experiments.

To explore the above hypothesis, two series of human experiments were conducted in two different task domains: process control and artificial grammar learning. Both tasks involve learning to control a system that operates according to complex, difficult-to-learn rules.

The process control task involved learning to control the temperature of a simulated nuclear reactor by controlling the number of fuel pellets. The appropriate number of pellets to use depended on the current temperature of the reactor, sometimes creating counterintuitive situations where increasing the number of fuel pellets decreases temperature. A noise element was included in the formula making the results somewhat uncertain over trials. This task is known to be difficult to learn and difficult to explain how one accomplishes the task when it is learned. In the process control experiments, we found that concurrent explicit reflection during practice either hindered learning the task or had no effect, even when solid hints were provided about what to look for while reflecting. The data seemed to suggest "just doing it" during practice was best, with some facilitation in learning through reflection after sessions of practice.

Then, to further validate our hypotheses from the process control domain concerning skill learning involving both implicit and explicit processes and their interaction, we extend our

studies to another domain---artificial grammar learning. In the artificial grammar experiments, participants learned to generate poison can labels on a computer simulation of a starship that had been invaded by enemy agents. Identification of poison food labels requires learning to identify sequences of letters generated by a finite state grammar. The findings in this domain confirmed our earlier hypotheses.

A computational cognitive architecture, CLARION, has been developed in this work to capture a range of data related to the interaction of implicit and explicit learning in the process control and artificial grammar domains. Simulation experiments have been carried out to further explore the interaction between implicit and explicit processes.

The work described in this report advances basic research in the areas of learning and cognition. One product of this effort is a conceptual framework, which addresses the ways these two types of knowledge interact to produce expertise. This framework (the CLARION cognitive architecture) suggests that human performance may be controlled by either a subconceptual knowledge base (the implicit mode) or application of a symbolic conceptual mental model (the explicit mode). Implicit control is fast but prone to error, particularly in early levels of skill acquisition. Explicit control is more accurate but slow to apply, and prone to loss by forgetting over a retention interval. We have found that reflection about how one is performing the task can be beneficial following periods of practice. However, it is often even more effective when learners are provided hints that direct their reflection in productive directions. These are important findings that advance our understanding of the interaction of the two types of knowledge. The computational cognitive architecture, CLARION, helps us to capture and explain (and eventually to predict) training and learning processes.

In the remainder of this report, the first section describes the initial simulation of process control data mentioned above. The next three sections report on the three human experiments on the process control task (as mentioned above). The five sections that follow report on the five human experiments on the artificial grammar task, as well as simulations. A general discussion section follows. Finally, a summary and conclusion section highlights a few important points and completes this report.

Modeling the Interaction of Explicit and Implicit Learning

In this section, a simulation examination of process control data is presented, which leads to some interesting hypotheses.

Specifically, the work reported in this section explicates the interaction between implicit and explicit learning processes in skill acquisition, contrary to the common tendency in the literature of studying each type of learning in isolation. It highlights the interaction between the two types of processes and its various effects on learning, including the synergy effect. This work advocates an integrated model of skill learning that takes into account both implicit and explicit processes; moreover, it embodies a bottom-up approach (first learning implicit knowledge and then explicit knowledge on its basis) towards skill learning. The simulation that follows shows that this approach accounts for various effects in the process control task data that have previously been reported in the literature (in addition to accounting for other human data, as described elsewhere). Notably, the simulation led to further human experimental work to be reported in the next eight sections, for testing and validating our ideas and hypotheses. The computational simulation generates these hypotheses concerning implicit vs. explicit processes, which necessitate validation with human experiments (to be described later in this report).

Introduction

The role of implicit learning in skill acquisition and the distinction between implicit and explicit learning have been widely recognized in recent years (see, e.g., Reber 1989, Stanley et al 1989, Willingham et al 1989, Proctor and Dutta 1995, Anderson 1993). Although implicit learning has been actively investigated, complex and multifaceted interaction between the implicit and the explicit and the importance of this interaction have not been universally recognized; to a large extent, such interactions have been downplayed or ignored, with only a few notable exceptions. Research has been focused on showing the lack of explicit learning in various learning settings (see especially Lewicki et al 1987) and on the controversies stemming from such claims. Similar oversight is also evident in computational simulation models of implicit learning (with few exceptions such as Cleeremans 1994).

Despite the lack of studies of interaction, there is increasing recognition that it is difficult, if not impossible, to find a situation in which only one type of learning is engaged (Reber 1989, Seger 1994, but see Lewicki et al 1987). Our review of existing data (see Sun et al 2001) has indicated that, while one can manipulate conditions to emphasize one or the other type, in most situations, both types of learning are involved, with varying amounts of contributions from each (see, e.g., Sun et al 2001; Stanley et al 1989, Willingham et al 1989).

Likewise, in the development of cognitive architectures (e.g., Rosenbloom et al 1993, Anderson 1993), the distinction between procedural and declarative knowledge has been proposed for a long time, and advocated or adopted by many in the field (see especially Anderson 1993). The distinction maps roughly onto the distinction between the explicit and implicit knowledge, because procedural knowledge is generally inaccessible while declarative knowledge is generally accessible and thus explicit. However, in work on cognitive architectures, focus has been almost exclusively on "top-down" models (that is, learning

first explicit knowledge and then implicit knowledge on the basis of the former), the bottom-up direction (that is, learning first implicit knowledge and then explicit knowledge, or learning both in parallel) has been largely ignored, paralleling and reflecting the related neglect of the interaction of explicit and implicit processes in the skill learning literature. However, there are a few scattered pieces of work that did demonstrate the parallel development of the two types of knowledge or the extraction of explicit knowledge from implicit knowledge (e.g., Rabinowitz and Goldberg 1995, Willingham et al 1989, Stanley et al 1989), contrary to usual top-down approaches in developing cognitive architectures.

Many issues arise with regard to the interaction between implicit and explicit processes: (1) How can we best capture implicit and explicit processes computationally? (2) How do the two types of knowledge develop along side each other and influence each other's development? (3) How is bottom-up learning possible and how can it be realized computationally? (4) How do the two types of knowledge interact during skilled performance and what is the impact of that interaction on performance? For example, the synergy of the two may be produced, as in Sun et al (2001). In the work described below, we will focus on the interaction and the synergy resulting from the interaction. The chief hypothesis of this work is the interaction between implicit and explicit knowledge is the key to understanding human skill learning.

An Integrative Model

Let us look into a model that incorporates both implicit and explicit processes.

Representation. The inaccessible nature of implicit knowledge may be captured by subsymbolic distributed representations provided by a backpropagation network (Rumelhart et al 1986). This is because representational units in a distributed representation are capable of accomplishing tasks but are subsymbolic and generally not individually meaningful (see Rumelhart et al 1986, Sun 1995); that is, they generally do not have an associated semantic label. This characteristic of distributed representation accords well with the inaccessibility of implicit knowledge. (However, it is generally not the case that distributed representations are not accessible at all but they are definitely less accessible, not as direct and immediate as localist representations. Distributed representations may be accessed through indirect, transformational processes.) In contrast, explicit knowledge may be captured in computational modeling by a symbolic or localist representations (Clark and Karmiloff-Smith 1993), in which each unit is easily interpretable and has a clear conceptual meaning, i.e., a semantic label. This characteristic captures the property of explicit knowledge being accessible and manipulable (Smolensky, 1988, Sun 1995). This radical difference in the representations of the two types of knowledge leads to a two-level model CLARION (which stands for Connectionist Learning with Adaptive Rule Induction ON-line; proposed in Sun 1997), whereby each level using one kind of representation captures one corresponding type of process (either implicit or explicit). Sun (1995, 1997, 2002), and Smolensky (1988) contain more theoretical arguments for such two-level models (which we will not get into here).

Learning. The learning of implicit action-centered knowledge at the bottom level can be done in a variety of ways consistent with the nature of distributed representations. In the

4

learning settings where correct input/output mappings are available, straight backpropagation (a supervised learning algorithm) can be used for the network (Rumelhart et al 1986). Such supervised learning procedures require the a priori determination of a uniquely correct output for each input. In the learning settings where there is no input/output mapping externally provided, reinforcement learning can be used (Watkins 1989), especially Q-learning (Watkins 1989) implemented using backpropagation networks. Such learning methods are cognitively justified: e.g., Shanks (1993) showed that human instrumental conditioning (a simple type of skill learning) was best captured by associative models (i.e., neural networks), when compared with a variety of rule-based models. Cleeremans (1997) argued that implicit learning could not be captured by symbolic models.

Specifically, $Q(x; a)$ is the "quality value" of action a in state x, output from a backpropagation network. Actions can be selected based on Q values, for example, using the Boltzmann distribution (Watkins 1989). We learn the Q value function through Q-learning, commonly used reinforcement learning algorithm (Watkins 1989). $Q(x; a)$ provides the error signal needed by the backpropagation algorithm and then backpropagation takes place (Rumelhart et al 1986).

The action-centered explicit knowledge at the top level can also be learned in a variety of ways in accordance with the localist representations used. Because of the representational characteristics, one-shot learning based on hypothesis testing (Nosofsky et al 1994, Sun 1997) is needed. With such learning, individuals explore the world, and dynamically acquire representations and modify them as needed, reflecting the dynamic (on-going) nature of skill learning (Sun, 1997; Sun et al 2001). The implicit knowledge already acquired in the bottom level can be utilized in learning explicit knowledge (through bottom-up learning; Sun et al 2001).

Initially, we hypothesize rules of a certain form to be tested (Dienes and Fahey 1995, Nosofsky et al 1994). When a measure of a rule (the IG measure) falls below the deletion threshold, we delete the rule. Whenever all the rules of a certain form are deleted, a new set of rules of a different form are hypothesized, and the cycle repeats itself. In hypothesizing rules, we progress from the simplest rule form to the most complex, in the order as shown in Figure 1, in accordance with those numerical relations used in human experiments (Berry and Broadbent 1988, Stanley et al 1989). (Other rule forms can be easily added to the hypothesis testing process. Since rules are tested in a parallel fashion, adding more rules will not drastically change the working of the model.)

$P = aW + b$
$P = aW1 + b$
$P = aW + cP1$
$P = aW1 + bP2$

Figure 1

The order of rules to be tested. $a = 1; 2$, $b = -1; -2; 0; 1; 2$, $c = -1; -2; 1; 2$, P is the desired system output level (the goal), W is the current input to the system (to be determined), W1 is the previous input to the system, P1 is the previous system output level (under W1), and P2 is the system output level at the time step before P1 .

The IG measure of a rule is calculated (in this process control task) based on the immediate reward at every step when the rule is applied. The inequality, $r >$ threshold, determines the positivity/negativity of a step and of the rule matching this step. Then, PM (positive match) and NM (negative match) counts of the matching rules are updated. IG is then calculated based on PM and NM (essentially as the positive match ratio).

The full CLARION model is highly comprehensive and therefore complex. The development of this cognitive architecture has taken many years of theoretical and experimental work. However, for the sake of maintaining a clear focus, only details most relevant to the simulations to be described below (a small subset of mechanisms) have been presented above. For further details of CLARION, see Sun (2002, 2003).

Simulation of Human Data

Simulation Focus. A number of well known skill learning tasks that involve both implicit and explicit processes were chosen to be simulated that span the spectrum ranging from simple reactive skills to more complex cognitive skills. The tasks include serial reaction time tasks, process control tasks, the Tower of Hanoi task, and the minefield navigation task. We focus on simulating process control tasks in this paper. We are especially interested in capturing the interaction of the two levels in the human data, whereby the respective contributions of the two levels are discernible through various experimental manipulations of learning settings that place differential emphases on the two levels. These data can be captured using the two-level interactive perspective.

We aim to capture (1) the verbalization effect, (2) the explicit (how-to) instruction effect, and (3) the explicit search effect. Through the simulations, it will be shown that the division of labor between, and the interaction of, the two levels is important.

To capture each individual manipulation, we do the following: (1) The explicit (how-to) instructions condition is modeled using the explicit encoding of the given knowledge at the top level (prior to training). (2) The verbalization condition (in which subjects are asked to explain their thinking while or between performing the task) is captured in simulation through changes in parameter values that encourage more top-level activities, consistent with the existing understanding of the effect of verbalization (that is, subjects become more explicit; Stanley et al 1989, Sun et al 1998). (3) The explicit search condition (in which subjects are told to perform an explicit search for regularities in stimuli) is captured

through relying more on the (increased) top-level rule learning, in correspondence with what we normally observe in subjects under the kind of instruction. (4) Many of these afore-enumerated manipulations lead to what we called the synergy effect between implicit and explicit processes: that is, the co-existence and interaction of the two types of processes leads to better performance than either one alone (Sun et al 2001). By modeling these manipulations, we at the same time capture the synergy effect as well.

General Model Setup. Many parameters in the model were set uniformly as follows: Network weights were randomly initialized between -0.01 and 0.01. Percentage combination of the two levels (through a weighted sum) is used: that is, if the top level indicates that action a has an activation value $l(a)$ (which should be 0 or 1 as rules are binary) and the bottom level indicates that a has an activation value $q(a)$ (the Q-value), then the final outcome is $v(a) = w1 * l(a) + w2 * q(a)$. The combination weights of the two levels were set at $w1 = 0.2$ and $w2 = 0.8$. Stochastic decision making with the Boltzmann distribution (based on the weighted sums) is then performed to select an action out of all the possible actions. Other parameters include numbers of input, output, and hidden units, the external reward, the rule deletion threshold, the backpropagation learning rate, and the momentum. Most of these parameters were not free parameters, because they were set in an a priori manner (based on our previous work), and not varied to match the human data.

For modeling each of these manipulations, usually only one or a few parameter values are changed. These parameters are changed as follows. To capture the verbalization effect, we raise the rule deletion threshold at the top level. The hypothesis is that, as explained earlier, verbalization tends to increase top-level activities, especially rule learning activities. To capture the explicit search effect, we increase the weighting of the top level in addition to raising the rule deletion threshold. The hypothesis is that explicit search instructions tend to increase the reliance on top-level rule learning. To capture the explicit instruction effect, we simply wire up explicit a priori knowledge at the top level.

Below we will describe only two simulations to illustrate our main points. Many other simulations may be found in other publications of ours (e.g., Sun 2002).

Simulating Stanley et al. (1989)

The Task. Two versions of the process control task were used in Stanley et al (1989). In the "person" version, subjects were to interact with a computer simulated "person" whose behavior ranged from "very rude" to "loving" (over a total of 12 levels) and the task was to maintain the behavior at "very friendly" by controlling his/her own behavior (which could also range over the 12 levels, from "very rude" to "loving"). In the sugar production factory version, subjects were to interact with a simulated factory to maintain a particular production level (out of a total of 12 possible production levels), through adjusting the size of the workforce (which has 12 levels). In either case, the behavior of the simulated system was determined by $P = 2 * W - P1 + N$, where P was the current system output, P1 was the previous system output, W was the subjects' input to the system, and N was noise. Noise (N) was added to the output of the system, so that there was a chance of being up or down one level (a 33% chance respectively).

There were four groups of subjects. The control group was not given any explicit how-to

instruction and not asked to verbalize. The "original" group was required to verbalize: Subjects were asked to verbalize after each block of 10 trials. Other groups of subjects were given explicit instructions in various forms, for example, "memory training", in which a series of 12 correct input/output pairs was presented to subjects, or "simple rules", in which a simple heuristic rule ("always select the response level half way between the current production level and the target level") was given to subjects. The numbers of subjects varied across groups. 12 to 31 subjects were tested in each group. All the subjects were trained for 200 trials (20 blocks of 10 trials).

The Data. The exact target value plus/minus one level (that is, "friendly", "very friendly", or "affectionate") was considered on target. The mean scores (numbers of on-target responses) per trial block for all groups were calculated. Analysis showed the verbalization effect: The score for the original group was significantly higher than the control group (F (1, 73) = 5.20; p < 0.05). Analysis also showed the explicit instruction effect: The scores for the memory training group and for the simple rule group were also significantly higher than the control group. See Table 1.

Table 1
The human data for the process control task from Stanley et al (1989)

	Human Data	
	Sugar Task	Person Task
Control	1.97	2.85
Original	2.57	3.75
Memory Training	4.63	5.33
Simple Rule	5.91	4.00

The Model Setup. The model was set up as described earlier. We used 168 input units, 40 hidden units, and 12 output units. There were 7 groups of input units, each for a particular (past) time step, constituting a moving time window. Each group of input units contained 24 units, in which half of them encoded 12 system output levels and the other half encoded 12 system input levels at a particular step. The 12 output units indicated 12 levels of subjects' input to the system. The learning rate was 0.1. The momentum was 0.1.

The rule deletion threshold was set at 0.15 for simulating control subjects. To capture the verbalization condition, the rule deletion threshold was raised to 0.35 (to encourage more rule learning activities). To capture the explicit instruction conditions, in the "memory training" condition, each of the 12 examples was wired up at the top level as simple rules (in the form of P1 → W); in the "simple rule" condition, the simple rule (as described earlier) was wired up at the top level. A reward of 1 was given when the system output was within the target range. In simulating the person task (a common, everyday task), we used pre-training of 10 blocks before data collection, to capture prior knowledge subjects likely had in this type of task.

The match. Our simulation captured the verbalization effect in the human data well. See Table1 and 2. We used a t-test to compare the "original" group with the control group in the model data, which showed a significant improvement of the original group over the

control group (p < :01), the same as the human data.

Table 2
The model data for the process control task from Stanley et al (1989)

	Model Data	
	Sugar Task	Person Task
Control	2.276	2.610
Original	2.952	4.187
Memory Training	4.089	5.425
Simple Rule	4.073	5.073

Our simulation also captured the explicit instruction effect, as shown in Table 2. We used pair-wise t-tests to compare the "memory training" and "simple rule" groups with the control group in the model data, which showed significant improvements of these two groups over the control group, respectively (p < :01).

Both effects point to the positive role of the top level. When the top level is enhanced, either through verbalization or through externally given explicit instructions, performance is improved, although such improvement is not universal (Sun et al 2001). They both showed synergy between the top-level explicit processes and the bottom-level implicit processes.

Simulating Berry and Broadbent (1988)

The Task. The task was similar to the computer "person" task in Stanley et al (1989). Subjects were to interact with a computer simulated "person" whose behavior ranged from "very rude" to "loving" and the task was to maintain the behavior at "very friendly" by controlling his/her own behavior (which could also range from "very rude" to "loving"). In the salient version of the task, the behavior of the computer "person" was determined by the immediately preceding input of the subject: It was usually two levels lower than the input $(P = W - 2 + N)$. In the non-salient version, it was determined by the input before that and was again two levels lower than that input $(P = W1 - 2 + N)$. Noise (N) was added to the output of the computer "person" so that there was a chance of being up or down one level (a 33% chance respectively).

Four groups of subjects were used: salient experimental, salient control, non-salient experimental, and non-salient control. The experimental groups were given explicit search instructions after the first set of 20 trials, and after the second set of 20 trials were given explicit instructions in the form of indicating the relevant input that determined the computer responses (W or W1).. 12 subjects per group were tested.

9
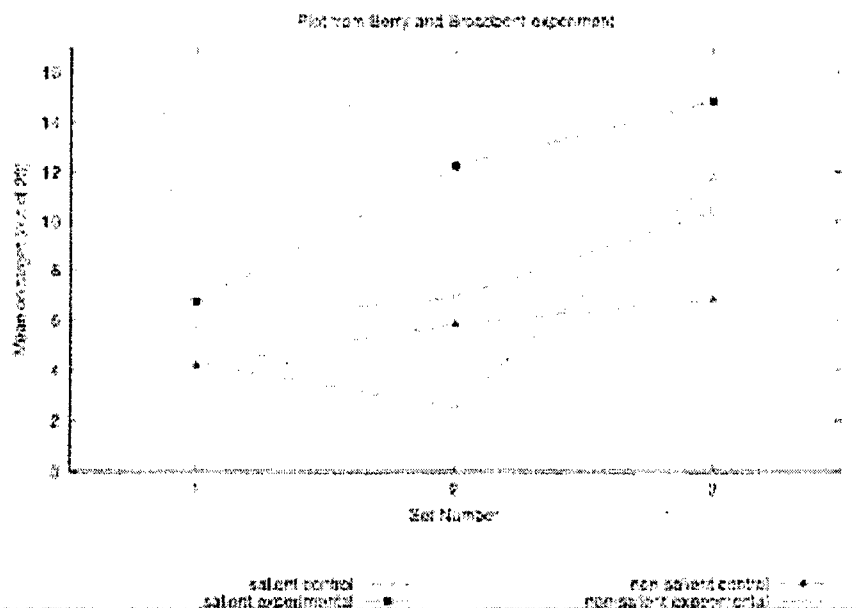
Figure 2
The data of Berry and Broadbent (1988)

The Data. The exact target value plus/minus one level (that is, "friendly", "very friendly", or "affectionate") was considered on target. The average number of trials on target was recorded for each subject for each set of 20 trials.

Figure 2 shows the data for the four groups of subjects for the three sets of trials. Analysis showed that on the first set, neither of the two experimental groups differed significantly from their respective control groups. However, on the second set, the salient experimental group scored significantly higher than the salient control group (p < 0.01), but the non-salient experimental group scored significantly less than the non-salient control group (p < 0.05). On the third set, both experimental groups scored significantly higher than their respective control groups (p < 0.01). The data clearly showed (1) the explicit search effect: improving performance in the salient condition and worsening performance in the non-salient condition; (2) the explicit instruction effect: improving performance in all conditions; as well as (3) the salience difference effect (during the 2nd set, under the explicit search condition).

The Model Setup. The model was set up similarly as described earlier for simulating Stanley et al (1989), except the following differences. The rule deletion threshold was set at 0.1 initially. To capture the explicit search effect (during the second training set), the rule deletion threshold was raised to 0.5 (for increased learning activities in the top level), and the weighting of the two levels was changed to 0.5/0.5 (for more reliance on the top level). To capture the explicit instructions given in this task (during the third training set), only rules that related the given critical variable to the system output were hypothesized and tested at the top level thereafter, in correspondence with the instructions (that is, $P = aW + b$, where W is the critical variable indicated by the instructions). The learning rate was 0.04. The momentum was 0.

10

The Match. We captured in our simulation of this task the following effects exhibited in the human data: the salience difference effect, the explicit search effect, and the explicit instruction effect. The results of the simulation are shown in Figure 3. On the first set, neither of the two experimental groups differed significantly from their respective control groups; however, on the second set, the salient experimental group scored slightly higher than the salient control group, but the non-salient experimental group scored slightly less than the non-salient control group. On the third set, both experimental groups scored significantly higher than their respective control groups ($p < 0.01$).

The data demonstrated clearly the explicit instruction effect (improving performance in all conditions), and showed to some extent the explicit search effect (improving performance in the salient condition and worsening performance in the non-salient condition), as well as the salience difference effect along with the explicit search effect. The data showed the extent and the limit of the synergy effect (in that the non-salient condition discouraged synergy
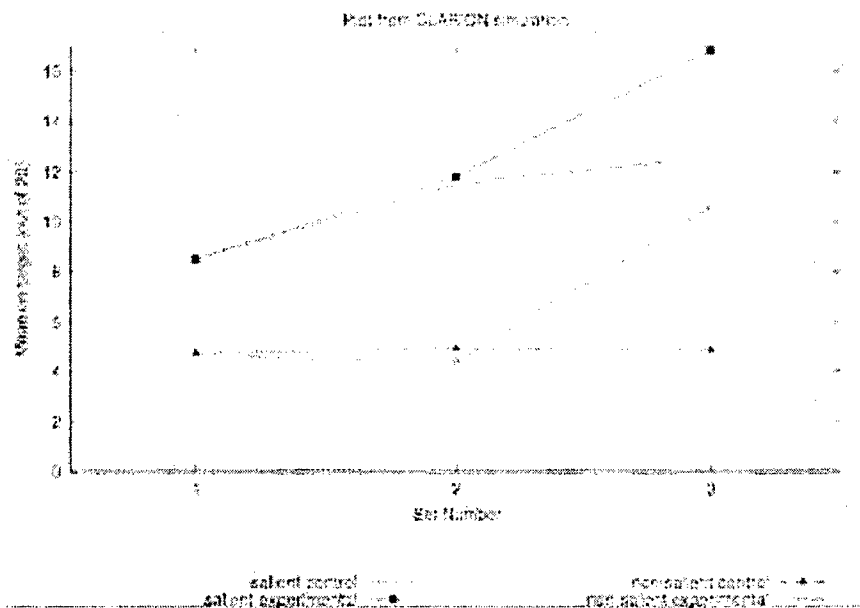
Figure 3
The simulation of Berry and Broadbent (1988)

Discussion

Although implicit learning is a controversial topic, the existence of implicit processes in skill learning is not in question. What is in question is their extent and importance. We allow for the possibility that both types of processes and both types of knowledge coexist and interact with each other to shape learning and performance, so we go beyond the controversies and the studies that focused mostly on the minute details of implicit learning (Gibson et al 1997).

The incorporation of both processes allows us to ask the question of how synergy is generated between the two separate, interacting components of the mind (the two types of

11

processes). The model may shed some light on this issue. Sun and Peterson (1998) did a thorough computational analysis of the source of the synergy between the two levels of CLARION in learning and in performance. The conclusion, based on the systematic analysis, was that the explanation of the synergy between the two levels rests on the following factors: (1) the complementary representations of the two levels: discrete vs. continuous; (2) the complementary learning processes: one-shot rule learning vs. gradual Q-value approximation; and (3) the bottom-up rule learning criterion used in CLARION. Due to space constraints, we will not repeat the analysis here. See Sun and Peterson (1998) for details. It is very likely, in view of the match between the model and human data as detailed in this paper, that the corresponding synergy in human performance results also from these same factors (in the main).

As a result of its distinct emphasis, CLARION is clearly distinguishable from existing unified theories/architectures of cognition, such as SOAR, ACT, and EPIC. For example, SOAR (Rosenbloom et al 1993) is different from CLARION, because SOAR makes no distinction between explicit and implicit learning, and is based on specialization, using only symbolic forms of knowledge. EPIC does not make the distinction either although it includes sensory-motor processes. Although ACT (Anderson 1993) makes the distinction, it is different from CLARION because traditionally it focuses mainly on top-down learning (from declarative to procedural knowledge).

The work reported thus far highlights the importance of the interaction of implicit and explicit processes in skill learning. It captures the interaction through a model that includes both types of processes. This modeling work reveals something new in the existing data (cf. Gibson et al 1997, Lebiere et al 1998). The contribution of this model lies in capturing human data in skill learning through the interaction of the two types of processes, and also in demonstrating the computational feasibility and psychological plausibility of bottom-up learning (Sun et al 2001). Note that many other simulations have been carried out that likewise show that the interaction between implicit and explicit knowledge during skill learning (see, e.g., Sun 2002 for details).

Now the question is how we verify the chief hypothesis of this model: The interaction between implicit and explicit knowledge is the key to understanding human skill learning. In the remainder of this report, we will address this question.

12

A great deal of research has been conducted over the last two decades to differentiate between implicit and explicit learning. Explicit learning is effortful (Norman, 1993) and results in a consciously available knowledge that can be readily verbalized. Implicit learning is like pattern recognition. For example, when we recognize a person's face, we are consciously aware of whom it is, but we have little conscious insight into what features were used to recognize the person. Implicitly acquired knowledge only tells us what to do; it does not provide a readily verbalizable set of rules to explain our behavior (e.g., Reber, 1967).

Several findings of this body of research suggest limited usefulness of implicit knowledge to support complex skills. Some studies (e.g., Dienes & Berry, 1997) provide evidence that implicit knowledge is so tied to specific training stimuli that it does not generalize beyond the exact instances experienced during training. Other research suggests that implicit knowledge is fragmentary and incomplete (e.g., Dulany, Carlson, & Dewey, 1984; Perruchet & Pacteau, 1990). In addition, research suggests that people have little confidence in implicitly acquired knowledge. They often think they are just guessing when applying their implicit knowledge (Chan, 1992; Dienes & Berry, 1997).

However, Mathews (1997) argued that these apparent limiting characteristics of implicit knowledge might be an artifact of the paradigms used to study it. Natural situations that depend heavily on implicit knowledge (natural language processing or pattern recognition) require extensive practice. Such tasks demand high levels of speed, accuracy and flexibility. Mathews (1997) suggested that experiments on implicit knowledge have focused too much on simple tasks because researchers were seeking cases of pure implicit (completely unconscious) knowledge. Typical experiments involve practice for less than 30 minutes. This amount of practice may be inadequate to develop levels of implicit knowledge that enable accurate and flexible utilization. Also, most real world situations do not involve pure implicit or pure explicit knowledge, but instead some blend of the two. Thus, it is important to study ways in which these two types of knowledge interact to influence performance on complex tasks (Sun, Merrill, & Peterson, 2001).

The impact of explicit reflection upon one's knowledge and thinking can vary. Facilitative effects of reflection have been found (Ahlum-Heath & DiVesta, 1986; Berry, 1983; Chi, Bassock, Lewis, Reimann, & Glaser, 1989; Chi, DeLeeuw, Chiu, & LaVarcher, 1994), however, reflection is not universally helpful. For example, verbalizing one's thoughts about difficult-to-verbalize aspects of one's knowledge can impair performance (*The verbal overshadowing effect*, Schooler & Engstler-Schooler, 1990). Indeed verbalization has been demonstrated to impair insight problem solving (Schooler, Ohlsson, & Brooks, 1993), analogy retrieval (Lane & Schooler, in press), affective decision-making (Wilson & Schooler, 1991), and memory for faces (see Meissner & Brigham, 2001 for a meta-analysis). In addition, when learning to perform complex tasks, learners may acquire invalid reflective knowledge in the form of mental models or verbalizable rules that lead to less than optimal performance (Reber, 1976; Reber, Kassin, Lewis, & Cantor, 1980). In short, the nature of the task (and we will argue, the nature of the reflection), can determine whether reflection has a positive, negative, or negligible impact.

The effect of different types of explicit reflection on process control and related (e.g., artificial grammar) tasks has been studied. One form of reflection involves simply instructing

13

participants to attempt to figure out the rules governing the behavior of the task. The effect of this manipulation has been mixed, sometimes decreasing the level of learning (e.g., Berry & Broadbent, 1988; Howard & Ballas, 1980; Reber, 1976; Reber et al., 1980) to having no effect (Dienes, Broadbent, & Berry, 1991; Dulany, Carlson, & Dewey, 1984), or improving learning (Berry & Broadbent, 1988; Reber et al., 1980). The primary mediating variable appears to be the salience of the rules governing the task. When the rules governing relations among the stimuli are salient or easy to discover, rule-search instructions can have a positive effect on learning (Mathews et al., 1989; Lee, 1995; Reber et al., 1980). However, rule-search instructions do not always facilitate performance in implicit learning tasks (Dulany et al., 1984; Lee, 1995; Mathews et al., 1989). In learning tasks involving rules that are extremely difficult to find (such as in the process control task), participants are likely to fall back on an implicit or memory-based mode to guide their performance (e.g., Mathews, et al., 1989).

Berry & Broadbent (1984) used the process control task to discover if verbal instruction on how to reach the target would affect task performance and verbalizable (explicit) knowledge similarly. They found that verbal instruction improved their participants' ability to control sugar production, except when combined with a requirement to verbally justify each response. Roussel (1999) investigated the effects of explicit reflection using the process control task by exposing learners to others' ideas about the task (other participants' policies or an experimenter-provided task hint), and by giving them the opportunity to discuss those ideas with other learners. Roussel's results demonstrated that certain types of explicit reflection can sometimes actually harm knowledge acquisition in this type of task. Assisted reflective practice, which involved a computer program designed to assist learners in thinking about their policies for controlling sugar production and to help them evaluate their policies by using them to perform the task, was found to be quite damaging to learning and performance. Another method for eliciting within-task reflection during task performance was to require participants to predict the outcome workforce size. As with assisted reflective practice, the participants had poorer task performance than did participants in a (non-prediction) control condition. Even the simplest method of reflection, involving giving learners pencil and paper along with instructions to use them to help them learn the task found no effect on learning. The present research investigated differences in interference effects on learning by varying the context of the task (Experiment 1), using occasional rather than continuous concurrent reflection (Experiment 2), and using the more casual form of concurrent reflection of taking notes during practice (Experiment 3). Post-task reflection was also examined.

Introduction

Experiment 1 replicates the findings of Roussel (1999), showing that explicit reflective practice interferes with learning to control sugar production in a process control task. Roussel proposed two mechanisms for the interference effect of reflection on performance. One was the generation of inaccurate explicit rules based on attempted reflection about the task. The second was interference with the implicit learning process (e.g., reflection acts like a secondary task). This experiment examined this interference effect in two different problem contexts: a sugar factory (replicating Roussel 1999) and in the context of controlling temperature in a nuclear reactor. The reactor control version of the task employed the exact same formula to control

14

output. However, the output variable is labeled reactor temperature (instead of sugar) and the input variable is labeled number of fuel pellets (instead of workers).

If the major negative impact of explicit reflection occurs because of generating inaccurate rules, we would expect a stronger interference effect of reflective practice in the original sugar production version. This is because the sugar factory version of the task offers a richer domain for generating overly general or inaccurate rules. Our participants are familiar with many things that could increase or decrease production in a factory (e.g., overcrowding, worker fatigue, firing less productive workers). Thus, when counterintuitive events happen in the sugar production version of the task, participants have a richer domain to draw on to generate rules. On the other hand, the reactor control scenario is relatively foreign and it seems mechanical. It would be difficult to think of complex but reasonable rules to account for the counterintuitive behavior of the system with this version of the task. Therefore, it was hypothesized that we would see a bigger interference effect from reflective practice in the sugar versus in the reactor control task. However, if Roussel's second factor, interference with the implicit learning process, is more important we might expect similar levels of reflective interference in both versions of the task.

Method

Participants and Design. Eighty six undergraduate students enrolled in introductory psychology courses at Louisiana State University were recruited to voluntarily participate in return for extra-credit. The experiment was arranged as a factorial design comprising three factors: task version (reactor control vs. sugar production) practice mode (reflective practice vs. experiential practice), and session (one through three).

The two primary dependent variables were performance, as indicated by the average unsigned deviation from target production during the test phase, and quality of the final policy. Policy quality was measured by using the policy to simulate performance of the sugar production task. The average unsigned deviation from target production achieved by the simulated policy was taken to be the policy quality. The procedure for evaluating policy quality will be described in detail below.

Process Control Task. One version of the process control task (Berry & Broadbent, 1984) used in this research has subjects imagine they are controlling a factory that produces sugar. The goal is to obtain a given target level of production (6,000 tons) on each trial. The subjects control a single variable, the number of workers employed at the factory. Production is affected by the number of workers in the following way: $P = (2 \times W) - P1 + N$. In this equation, P = current sugar production, W = number of workers input by subject, $P1$ = previous level of sugar production, and N = noise (a random element).

This research compared two versions of this task, the sugar production version and the reactor control version of the control task. The reactor control task was exactly the same as the sugar production task in all aspects except the cover story and labels of the input and output variables. The task was described as a simulated nuclear reactor. Their task was to maintain the reactor temperature as close as possible to the target level (6,000 degrees). On each trial they had to input a new input level for number of fuel pellets.

Task trials were grouped into blocks of ten trials and each block began with a randomly selected production level. Figure 4 shows the graphical display seen by participants in the

reactor control version of the task. The graph on the left side of the screen plots input responses over trials and the graph on the right plots output levels across trials. On the left graph, the number of fuel pellets entered on each trial is displayed on the horizontal axis. On the right graph, reactor temperature level is represented on the vertical axis of the graph. The dashed horizontal line shows the target temperature level. The horizontal axis represents the sequence of trials. Each trial output is represented by an 'X' on the graph. At the end of each block, the display was cleared and a new graph displayed for the next block of trials. Temperature (sugar production in the sugar factory version) was allowed to vary from 1000 degrees to 12000 degrees. Participants were allowed to select an input value (for fuel pellets or workers) ranging from 100 to 1200 in multiples of 100. The target production was fixed at 6000 tons. The only difference in the sugar production version of the task was the labels associated with input and output variables.
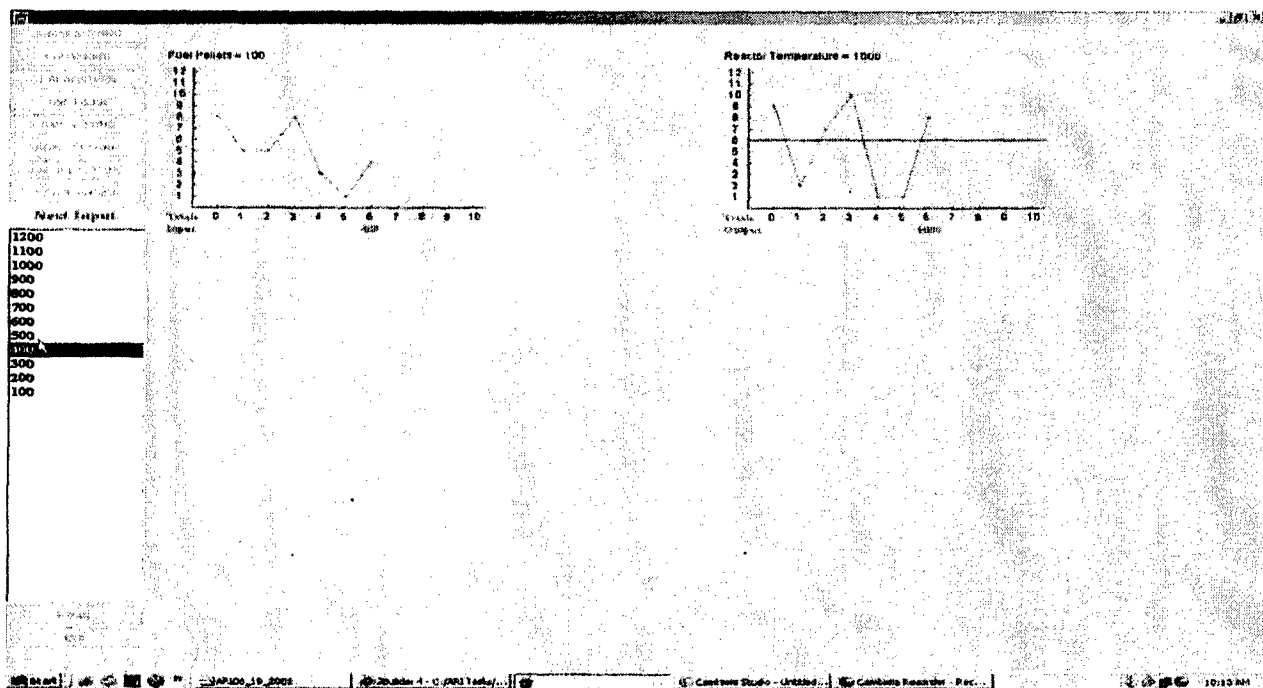
Figure 4.
Graphical display seen by a participant performing the reactor control task on the sixth trial in a block of 10 trials.

The relationship between number of workers and sugar production was identical to that used by Roussel (1999). The main dependent measure was the mean unsigned deviation from target production, in tons, across a block of ten trials. Because the target production level was always 6000 tons, the dependent measure could vary from a minimum of zero, if on target for every trial, to a maximum of 6000 tons away from target level. Chance performance was defined as the mean unsigned deviation that would be achieved by entering a random value for workers on every trial. Chance performance was thus determined to be 4206 tons. Best performance

16

possible is about 600 tons off target on the average because of the noise element in the task control equation.

Procedure. Participants were tested in groups ranging from three to five individuals. Each group was randomly assigned to one of the four conditions. Regardless of condition, all participants completed three sessions, one per day. For all participants, the three sessions were completed within seven days. All participants performed 20 minutes of practice followed by 10 blocks (100 trials) of test. Additionally, participants in the reflective practice condition had up to 15 minutes at the end of each session to write (Session 1) or revise (Session 2-3) their policy on how to control the task.

First Session. In the first session, all participants were told that they were to take on the role of manager of either a simulated sugar production factory or a simulated nuclear reactor. They were informed that their job was to learn how to achieve and maintain a target level of output by interacting with the simulation. They were further informed that the only variable they could control was the one input variable (either workers or number of fuel pellets). Thus, their task was to learn the relationship between workforce size and production level in the simulated sugar factory conditions, and amount of fuel and reactor temperature in the simulated nuclear reactor conditions. Participants in the reflective practice conditions were also told that they would be required to write a policy or set of instructions for someone else to perform the task at the end of each session.

After receiving instructions, all participants were given 20 minutes to practice or interact with the simulation program. In Session 1 all participants simply performed the process control task at their own pace during the practice period.

The test comprised ten blocks of 10 trials of the same task. The participants were allowed up to 30 minutes to complete the test. The participants were informed that their goal was to stay as close as possible to the target production level and that there would be a $50 reward for the best performance.

After completing the test, each participant in the reflective practice condition was given 15 minutes to write down his or her policy for controlling sugar production or the nuclear reactor. These participants were told that someone else would try to perform the process control task using only the instructions they provide. They were also told there would be an additional $50 reward for the best policy, determined by the best performance using a participant's policy to perform the task. Participants were allowed up to 15 minutes to write their policy. They were asked to write each statement of their policy on a numbered page, giving each statement a new number.

Second and Third Sessions. In all reflective practice conditions participants were returned their written policies from the previous session. The same practice-test-write policy sequence used in the first session was followed for the second and third sessions. However, before beginning to practice, all reflective practice participants were told that they would have to write a new policy at the end of the session and therefore, they should be thinking about how to improve their policy as they practiced.

Participants in the experiential practice conditions simply performed the process control task at their own pace during the practice period, as they did in Session 1. Participants in the reflective practice conditions were required to record on a log sheet for every trial which particular statement of their written policy they were following, the number of workers to be

17

used according to their rule, and the production level they expected to achieve. After entering this information on their log sheet they could type in their selected input level and the computer calculated and displayed the new production level.

The test portion of the second and third sessions was the same as in Session 1. It was a 10-block sequence of the process control task. However, this time, reflective practice participants were allowed to refer to their written policies from the previous session as they performed the test. Participants in the reflective practice conditions were not required to use the log sheet during the test.

At the end of the session, participants were instructed to write new policies based on the performance of their old policies. They were informed that they could include any part or all of their old policies in the new one. After the end of the third session, all participants were debriefed and given a slip for their extra credit points.

Policy Evaluation. The ratings were determined by using them to perform the sugar production task for 10 blocks of trials. On each trial, a rater selected the most appropriate rule from the policy and entered the indicated number of workers. The most appropriate rule was considered to be the one that matched the current situation and was the most specific in its range of application. For example, consider the following two rules: (1) "If you are above the target production of 6000 then you should decrease the size of the workforce"; and (2) "If current production level is between 8000 and 10000 tons then you should use 800 workers." Both rules would be applicable to any trial on which current production level is 9000 tons. However, the second rule is more specific (i.e., applicable in fewer situations) and would be chosen by the rater. On trials where no rule applied, the rater entered the same number of workers used on the previous trial, unless it was the first trial. In this situation, the rater entered a randomly selected number of workers. On trials where the policy indicated only a range of workers (e.g., more workers, or a high number of workers) the following actions were taken: (a) "more, or less, workers than X" was interpreted as a randomly selected value of workers between X and the maximum or minimum number of workers allowed, respectively; (b) "a high, or low, number of workers" was taken to mean a randomly selected number of workers above 750 or below 450 respectively; and (c) "an increasing, or decreasing, number of workers" was interpreted the same as in (a). A random number generator (computer program) assisted the rater in selecting random values.

Results

The mean deviation from target level as a function of session, task, and practice mode is presented in Table 3. Since the reflective practice manipulation was not implemented until Sessions 2 and 3, these data were analyzed separately from Session 1. In Session 1 the practice mode conditions only differed in that participants in the reflective practice conditions were told that they would write a policy at the end of the session. This knowledge might have stimulated more reflective thinking during practice in Session 1 in the reflective practice groups.

Table 3.
Means and Standard Error (in Parentheses) of Deviation from Target Level on Test as a Function of Session, Task, and Type of Practice.

Sugar Production		N	Session 1	Session 2	Session 3
	Reflective Practice	19	2919 (144)	2612 (192)	2424 (193)
	Experiential Practice	22	2712 (134)	2504 (141)	2256 (164)
Reactor Control					
	Reflective Practice	23	2837 (131)	2730 (116)	2476 (129)
	Experiential Practice	20	2732 (141)	2105 (163)	1752 (129)

Total Research Trials. As in the original Roussel (1999) research, the reflective and experiential practice conditions were equated in terms of practice time. However, the reflective practice groups performed the task at a much slower rate in order to reflect on applying their policy and logging their choices and results. Thus, by the end of Session 3, the experiential practice conditions had performed a lot more trials of the task. The mean number of total research trials across the four groups were: 1832 trials in the experiential, reactor control group, 1949 in the experiential, sugar production group, 734 in the reflective practice, reactor control group, and 750 in the reflective practice, sugar production group.

Session 1 Performance. There were no significant differences between any of the groups in Session 1. Apparently informing participants in the reflective practice conditions that they would be required to write a policy at the end of the session did not affect performance.

Session 2-3 Performance. Performance means for Sessions 2 and 3 were analyzed using a repeated measures ANCOVA. The three factors included in the ANCOVA were session, practice mode, and task. Session was the repeated measure factor. Total research trials was the covariate. There was a significant effect of total research trials, $F(1,78) = 6.711$, $p<.05$. Performance improved across sessions, $F(1, 78) = 6.325$, $MSE = 90661$, $p < .05$, indicating that participants were learning to control task output. Participants in experiential practice conditions consistently outperformed their reflective practice counterparts, $F(1,78) = 14.706$, $MSE = 831828$, $p < .01$, replicating the Roussel (1999) finding of a negative effect of reflective practice. There were no other significant effects or interactions. It should also be noted that the variability across participants tended to be higher in the sugar task with reflective practice (see standard error values in Table 1).

The Effect of Assisted Reflective Practice on Reflective Knowledge. The mean simulated performance of the final session policies was 3315 in the reactor control version of the task and 3048 in the sugar production. These means are not significantly different indicating that policy quality did not differ as a function of task version. The mean correlation between policy quality and final test performance was .38 in the sugar version of the task and .57 in the reactor control version. Only the correlation for the reactor control version of the task was significant. Thus,

19

there is more evidence of explicit knowledge in the reactor control version of the task.

Discussion

This experiment replicated the negative effect of explicit reflective practice found by Roussel (1999) in two versions of the process control task. However, the prediction that a larger interference effect would occur in the more familiar sugar factory version of the task was not supported. This result suggests that richness of potential rules in a domain is not related to size of the reflective practice interference effect. Perhaps the large quantity of overly general or inaccurate rules found by Roussel (1999) was more directly linked to the group discussions in their experiments rather than assisted reflective practice. Or, alternatively, perhaps the large set of "bad" rules was a by-product of poorer implicit learning rather than a cause of poor performance on the task. However, the negative effect of reflective practice was replicated in both versions of the task in our experiment. Therefore, interference with the implicit learning process rather than generating overly general rules while reflecting seems to be the major cause of the interference effect of reflection on task performance. Experiment 2 further tests this notion by implementing a partial reflective practice condition.

Introduction

Roussel (1999) suggested one interpretation of the negative effect of reflective practice was that participants generated and used overly general or incorrect explicit rules about the process control task. However, Experiment 1 showed that richness of domain knowledge (factory vs nuclear reactor) did not alter the negative effect. Experiment 2 uses a partial reflective practice condition to see if occasional rather than continuous concurrent reflection disrupts learning of the task. If concurrent reflection leads participants to generate bad rules as suggested by Roussel (1999), then having participants reflect on even a small subset of trials would still lead to disruption. Participants would generate bad rules on the reflection trials and continue to use these rules on subsequent trials. On the other hand, if only *continuous* concurrent reflection disrupts implicit learning, then the source of the interference might be interference with the implicit learning process. Participants who have partial reflective practice could still learn the task implicitly without interference on the non-reflection trials.

Method

Only the reactor control version of the task was used in the remaining experiments in this research. Experiment 2 employed the same reflective practice procedure used in Experiment 1. However, rather than a fixed amount of practice time, a fixed number of trials was used to equate the amount of practice trials between the partial reflection and the experiential conditions. This was done to insure that any disruptive effect of partial reflection could not be attributed to fewer practice trials resulting from the slow reflective process (even though the covariate analyses in Experiment 1 suggested this was not the case). In each session all participants practiced for 30 blocks of trials, then they took a test consisting of 30 blocks. Reflective practice participants wrote a policy at the end of each session. In Sessions 2-3 **partial reflective practice** participants used their previous session policy to perform reflective practice on the first 10 blocks of practice. The second 20 blocks were performed without reflective practice. **Experiential practice** participants simply practiced 30 blocks of the task each session and they did not write policies at the end of each session. All other aspects of the design were identical to Experiment 1. There were 18 participants in each of the two conditions.

Results

Means and standard error for test performance in all three sessions are shown in Table 4. An ANOVA on Session 1 (before the reflective practice procedure was implemented) showed no significant difference between groups on performance. An ANOVA on Sessions 2-3 revealed only a significant effect of Session, $F(1,34)=5.927$, $\underline{MSE} = 64388$, $p<.05$. There was no effect of reflection.

Table 4
Means and Standard Error (in Parentheses) for Test Performance

	Session 1	Session 2	Session 3
No reflection	2653 (150)	2509 (174)	2421 (187)
Partial Reflective Practice	2398 (146)	2435 (174)	2233 (188)

Final policy quality in the partial reflective practice group was 3065. The correlation between policy quality and final test performance was significant (r = .54), indicating there was some level of valid knowledge in the policies.

Discussion

Clearly, just activating explicit thinking during practice was not sufficient to produce the negative effect of reflection. This finding does not support the Roussel (1999) interpretation of the effect in terms of generating bad explicit rules. Rather, the results suggest that the negative effect found in Experiment 1 may have resulted from interference of the reflective practice procedure with implicit learning processes. Perhaps this procedure interrupts the process of storing information about experiences in the memory used to support implicit knowledge. If this latter interpretation is correct, a less structured form of reflection might not interfere with implicit learning and, instead, facilitate task performance.

Process Control Experiment 3

Introduction

Experiment 3 uses a speeded version of the task to block concurrent reflection, varied the opportunity for **post** task reflection, and used a powerful set of hints thought to enhance learning of this task (see Roussel 1999). The hints consisted of three examples of good rules to apply when current output was each of three specified levels. For example, If current temperature is 4000 tons, then use 500 fuel pellets. Roussel found that both performance and policy quality were enhanced when learners were provided with four such examples combined with a general statement that said: "The number of workers should always follow the level of production. That is, when production is high, you need a lot of workers and when production is low, you need few workers. Similarly, when production is near the middle, you should use a moderate level of workers, not high and not low."

Roussel also found that providing the rule exemplars with the general statement or just the general statement alone, both enhanced learning. However, they did not provide example rules alone, so we can not be sure the example rules would be effective by themselves. Logically, however, they should be. Good policies are generally lists of just such specific rules. The learner would simply have to fill in the rest of the 12 mini rules when she discovers them during practice. However, this filling in of a look up table would seem to require conscious effort in the form of reflection either during or after practice.

To facilitate this type of reflection during practice, some participants were provided with pen and paper and they were encouraged to take notes whenever they wished. Roussel found this type of informal task reflection during practice did not facilitate or impair learning. Participants allowed such informal concurrent reflection used the regular self-paced version of the reactor control task. Participants not allowed concurrent reflection used a fast paced (5 sec per trial) version of the task designed to minimize concurrent reflection during practice.

After each 15 min session of practice all participants performed another task for five minutes. For participants allowed post task reflection, this task consisted of writing a policy about how to perform the reactor control task. For participants not allowed post task reflection, this interim task consisted of watching and rating video advertisements.

It was expected that some reflection would be necessary to benefit from the hints. We also predicted that post task reflection would be the most effective because it would not interfere with the implicit learning processes during training.

Method

Experiment 3 was a 2 X 2 X 2 X2 factorial design with three between participant factors, concurrent reflection (or not), post task reflection (or not), and hint (or not) and one within participant factor (Session 1 and 2). As in Experiment 1, timed periods of practice were used rather than set numbers of trials (as in Experiment 2). Each session consisted of two sequences of 15 minutes of practice followed by five minutes of policy writing or advertisement rating. Participants who did not write a policy were asked to rate five videotaped commercials per five

minute period following each sequence. After the second five minutes of policy writing or ad evaluation, 10 blocks of test were administered.

Post task reflection consisted of five minutes of writing a policy for controlling the reactor following each set of practice trials. Participants that were not allowed post task reflection performed a distracter task, rating video advertisements, during that five-minute interval. The ads were rated using a five point Likert scale for effectiveness in selling the product.

Concurrent reflective practice participants were encouraged to take notes during practice. They were also allowed to refer to their notes and/or policies during the test.

The hint consisted of providing three examples of good rules for specific output levels. The hint was:

If current temperature is 1000 then use 400 pellets
If current temperature is 4000 then use 500 pellets
If current temperature is 7000 then use 700 pellets

Results

The means for test performance are presented in Table 5. The only significant effects in the ANOVA on test performance were: hint, $F(1,197)=6.86$, MSE=950556, $p<.01$, post task reflection, $F(1,197)=3.80$, MSE=950556, $p=.05$, Test, $F(1,197) = 163.32$, MSE=243090, $p<.01$, and the test by post interaction, $F(1,197) = 9.91$, MSE=243090, $p<.01$. Thus, the exemplar hint helped performance even without any opportunity for reflection (compare the top two rows in Table 3). Casual concurrent reflection was neither damaging nor helpful to performance. Post task reflection was beneficial (with or without the hint), but it only enhanced learning in the first session.

Table 5.
Test Performance Means and Standard Error (in Parentheses) as a Function of Reflection and Hint

No Reflection	Test 1		Test 2	
Exemplar Hint	2266	(149)	1463	(137)
No Hint	2360	(154)	1572	(143)
Concurrent Only				
Exemplar Hint	2474	(144)	1522	(133)
No Hint	2662	(180)	2077	(166)
Post Only				
Exemplar Hint	1874	(157)	1402	(145)
No Hint	2310	(168)	1791	(154)
Concurrent Plus Post				
Exemplar Hint	2060	(164)	1561	(151)
No Hint	2143	(161)	1741	(148)

Since participants in the no-concurrent reflection group had to respond very fast, they would have experienced more practice trials than the concurrent reflection groups. The mean number of practice trials across both sessions in these groups were: 2571 in the no reflection group, 2540 in the post only reflection group, 1774 in the concurrent only reflection group, and 1566 in the concurrent and post reflection group.

An analysis using final test performance as the dependent variable and total practice trials as a covariate showed a significant effect of practice trials, $F(1,195) = 35.664$, MSE=465851, $p<.001$, hint, $F(1,195) = 4.168$, $p<.05$, and a strong negative effect of concurrent reflection, $F(1,195) = 23.735$, MSE=465851, $p<.001$. The hint by post task reflection by concurrent reflection was also significant, $F(1,195) = 4.843$, MSE=465851, $p<.05$. The adjusted means from this analysis are shown in Table 6

Table 6.
Means for Test in Session 2 Adjusted to Equate Total Practice = 2125 Trials.

No Reflection	Test 2
Exemplar Hint	1356 (128)
No Hint	1284 (140)
Concurrent Only	
Exemplar Hint	1671 (125)
No Hint	2240 (155)
Post Only	
Exemplar Hint	1285 (135)
No Hint	1541 (148)
Concurrent Plus Post	
Exemplar Hint	1891 (151)
No Hint	1938 (140)

The means for policy quality are presented in Table 7. The ANOVA on Policy quality indicated a significant effect of hint, $F(1,94) = 8.74$, $p<.01$ and policy order $F(3,282) = 28.89$, $p<.01$, MSE = 524765. Thus policy quality steadily improved across attempts and sessions and the exemplar hint increased policy quality. There were no other significant effects.

Table 7.
Means and Standard Error (in Parentheses) of Deviation from Target Level on Policy Quality as a Function of Session, Task, and Type of Practice.

Hint Reflection	Session 1 First Policy	Session 1 Second Policy	Session 2 First Policy	Session 2 Second Policy
Exemplar Hint Concurrent and Post	3389 (153)	3060 (197)	2467 (203)	2217 (224)
Exemplar Hint Post Only	3258 (147)	2660 (190)	2437 (195)	2540 (215)
No Hint Concurrent and Post	3576 (150)	3344 (193)	2859 (199)	2697 (219)
No Hint Post Only	3675 (156)	3264 (202)	3174 (208)	2916 (228)

Discussion

This experiment used a very innocuous form of concurrent reflection—just encouraging participants to take notes when they discover something new. One would expect that this type of reflection, especially when combined with the exemplar hints telling participants what to look for, would be very beneficial to learning. If the effect of reflection on performance resulted from hypothesis testing or in some way explicitly figuring out the rules of the game, such as finding the correct responses to fill in a look-up table (Dienes & Fahey, 1995), concurrent reflection with the hint should have been very helpful. Whenever a correct response was found it could be written down until all 12 possible correct responses were discovered. However, the ANOVA on test performance indicated no positive effect of concurrent reflection. In fact, while the difference in this analysis was not significant, the means are in the direction of an interference effect of concurrent reflection rather than a positive effect. Even more surprising, the ANCOVA, with total practice trials as the covariate, showed a very strong negative effect of casual concurrent reflection. That is, when equated for number of practice trials, the negative effect of casual reflection during practice gets stronger. Post task reflection was beneficial, but only early in learning. The strong message of these data is "just do it" is the way to learn this task. Don't think about it. Participants that neither reflected during (concurrent) nor after (post) practice ended up with the best scores in Session 2.

The most surprising result was the finding that even when participants were given virtually no time to reflect during practice (because of the speeded task) or after practice (the rating advertisement task filled this interval), the hint was just as effective. The hint was also effective in enhancing valid explicit knowledge of the task, as demonstrated by its effect on policy quality. Given that thinking about the task seems to have primarily negative consequences, how are we to explain the positive effect of the exemplar hint? We think the hint changes the way participants perceive the task. Perhaps it causes them to focus more on trial by trial changes in the relevant variables. Such a change in attention or encoding appears to enhance implicit learning of the task. We think the positive effect of hint on explicit knowledge (policy quality) results from bottom-up learning processes. In other words, the rules are

26

discovered implicitly but eventually become conscious and are transformed into the explicit rules used in the policies.

Below, we extend our studies to another domain---artificial grammar learning, in an effort to further validate our hypotheses from the process control domain concerning skill learning involving both implicit and explicit processes and their interaction.

Introduction

In the process control experiments we found that concurrent explicit reflection during practice either hindered learning the task or had no effect, even when solid hints were provided about what to look for while reflecting. The data seemed to suggest "just doing it" during practice was best, with some facilitation in learning through reflection after sessions of practice. Our goal in the following experiments was to examine the effects of similar training variables in another well studied implicit learning domain, artificial grammar experiments. We were also interested in examining these effects in situations that required both speed and accuracy of decisions. Therefore, we transformed the artificial grammar paradigm to a situations where participants had to react dynamically during the test to respond to cues provided by the computer (two letters in a valid string) and quickly generate a response (the rest of the string) that was close (70% correct) to a valid string. This task has some ecological validity to natural language learning in that a child need not be completely correct grammatically for a parent to understand and respond. Here too, our participants learning this artificial language needed only to approximate a valid string to be rewarded by the computer. Our test also removes potential valid responses from the set of possible strings as they are successfully generated, forcing the learner to encounter a wide range of possible valid strings. Thus, good learning of a few valid strings will not support good performance on the test.

Most theorists accept that some sort of implicit memory of experienced instances (either a neural network, a database of instances or sets of instance fragments) is the underlying basis for implicit knowledge (Knowlton & Squire, 1996; Manza & Reber, 1997; Mathews, 1991; Vokey & Brooks, 1992; Whittlesea & Dorken, 1993). However, there are still many questions about what type of training might be optimal for developing such an implicit memory bank of experienced instances.

Some researchers emphasize the storage of intact exemplars with performance based on the nearest neighbors in the memory bank (Brooks, 1978; Vokey & Brooks, 1992). Hence, a larger database of exemplars should be beneficial when comparing similarities between novel and stored exemplars (Whittlesea & Wright, 1997). Other researchers propose that this database contains partial memories of exemplars (Mathews, 1991), memories of chunks of exemplars (Servan-Scheiber & Anderson, 1990), or acquired knowledge of bigrams and trigrams and their frequencies (Perruchet & Pacteau, 1990). This partial memory view might depend more on the representativeness of experienced instances rather than having a large set of instances in memory.

Very little research has examined the effects of mixing implicit and explicit training. Reber, Kassin, Lewis, & Cantor (1980), in an experiment using a finite-state grammar, found that briefly exposing participants to the actual diagram of the grammar (explicit training) prior to

28

training with instances (implicit training) resulted in better performance on a string discrimination test. In contrast, Mathews et. al. (1989) found no advantage of mixed training with a finite-state grammar, but did find a beneficial effect of mixed training with a biconditional grammar. The explicit training task used in the latter research consisted of learning to correct invalid strings (the edit task). The implicit training task consisted of recognizing an exact copy of a valid string presented before each trial (match task). For participants learning the biconditional grammar, Mathews et. al. found that the group that had implicit training followed by explicit training performed better than all other groups.

The present series of experiments examines mixing training across sessions as well as an integrated type of training designed to provide simultaneous experience with exemplars (implicit training) and knowledge of the structure of the grammar (explicit training). This new training method is called exemplar diagramming (ED). In this training task, participants traced each training exemplar through a diagram of the artificial grammar. Thus, they processed exemplars (implicit learning) within the context of the grammar (explicit learning).

Method

Two training tasks were contrasted in Experiment 1. One training task, the exemplar processing or EP task, required participants to hold instances in memory long enough to copy them on a response sheet (see Panel A of Figure 5). The other task, exemplar diagramming or the ED task, required participants to trace the exemplars through a diagram of the grammar (see Panel B of Figure 5). This experiment also explored the effect of training set size. All groups had a set of 88 instances to process. However, the small training set consisted of 22 different exemplars repeated randomly four times while the large training set consisted of 88 different exemplars.

Performance was tested using the cued-generate test (Mathews & Cochran, 1998). This test requires generation of a large variety of exemplars based on minimal retrieval cues (two randomly selected letters). We expected that the explicit knowledge of the grammar obtained during the ED task would enhance performance. An explicit representation of the grammar could provide retrieval cues to help access relevant stored exemplars in the implicit memory bank. It could also enhance efficiency and accuracy of string generation by providing a means for correcting errors or omissions in memory traces of exemplars. Also, some researchers suggest that (purely) implicit knowledge is inflexible (Stadler, Justin, & Shana, 2000; Dienes & Altman, 1997). Thus, the purely implicit database created by exemplar processing in the EP task might function poorly in enabling generation of diverse sets of exemplars. Therefore, we expected the exemplar diagramming participants to outperform the exemplar processing participants on the cued-generate test in terms of efficiency (proportion of acceptable strings generated per attempt), and accuracy (number of perfect strings generated). However, using explicit knowledge is known to be a comparatively slow process (Reber et. al., 1980; Norman, 1993). Thus, the purely implicit (EP) group might respond faster. Overall achievement on the test (number of strings generated during a test session) might depend on the optimal balance of speed and accuracy given the task constraints (e.g., 20 minute time limit and 70% correct letter match required in generated strings).

Participants. Ninety-two undergraduate students taking a variety of psychology courses

at Louisiana State University participated in the experiment. All participants were volunteers and received extra credit for their participation.

    Materials. The finite-state grammar used by Mathews et. al. (1989) was used in this experiment (see Panel B of Figure 5). This grammar generates 177 exemplars ranging in length from 5 to 11 letters. Two representative subsets of exemplars from this grammar were used as training stimuli. One subset consisted of 88 exemplars and was termed the "large set". The other subset consisted of 22 exemplars and was termed the "small set". The exemplars in the small set were selected to illustrate all of the grammar paths and the effects of the two loops in the grammar (see Panel B of Figure 5). The small set was randomly repeated four times to make the number of instances equivalent in the two training sets. Thus, participants receiving either the large or small training set had a total of 88 instances available for their training task. Each exemplar from both training sets was typed onto labels and affixed to the center of a rolodex card. Both exemplar sets were presented randomly on cards bound to a rolodex base.

    Two response sheets were used for the different training tasks. The response sheet used by the exemplar processing (EP) groups consisted of six rows and twelve columns of circles. The six letters from the artificial grammar were printed vertically along the left side of the sheet. Along the top, the numbers one through twelve were printed horizontally, representing the serial order of the letters in an exemplar (see Panel A of Figure 5). The second response sheet was a transition diagram of the Mathews et al.'s (1989) artificial grammar used by the exemplar diagramming (ED) groups. It contained spaces to write the letters of the exemplars at the appropriate transition points within the grammar (see Panel B of Figure 5).
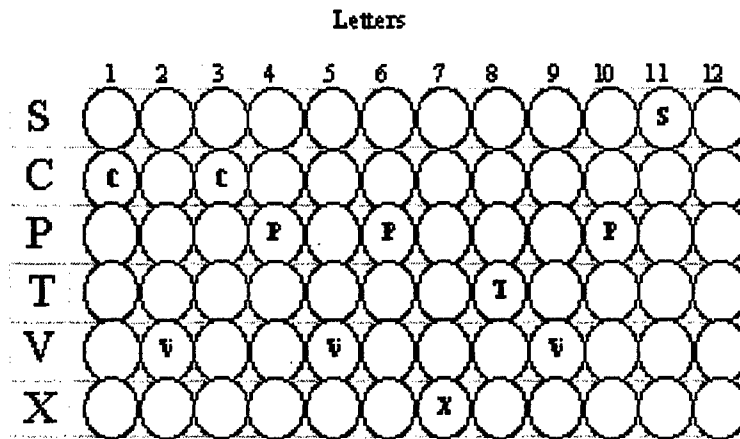
Panel A
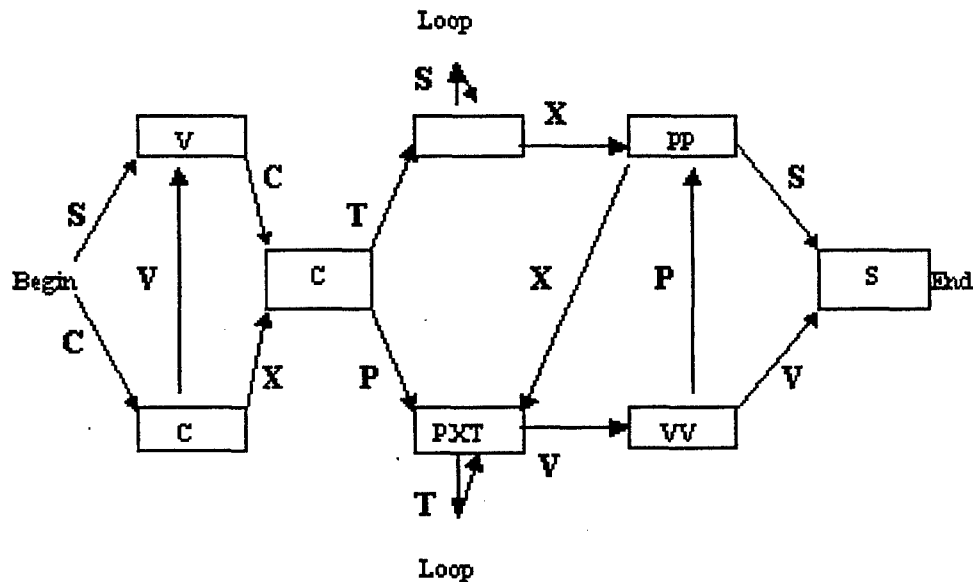
Panel B

Figure 5. *Training tasks. A. Bubble Sheet Diagram. The bubble sheet used by the participants to perform the exemplar processing (EP) training task. In the diagram the valid letter string CVCPVPXTVPS is inserted to illustrate the proper method used. B. Transition Map of the Grammar. The diagram used by participants to perform the exemplar diagramming (ED) training task. The same exemplar is traced through the map to illustrate proper insertion.*

31

Design. The design was a 2 x 2 x 3 (training task x research set length x session) factorial. The two training tasks (EP versus ED) and the length of the exemplar sets (large versus small) served as between-subjects factors. The three 1-hour weekly sessions served as the within-subjects factor. Twenty-three participants were randomly assigned to each of the four conditions.

Procedure. Participants were tested in groups of up to four. There were three 1-hour sessions scheduled one week apart. Each session began with a 20-minute training phase requiring participants to perform either the EP or ED training task. Each training phase was followed by a 20-minute cued-generate test.

As in Mathews, Roussel, Cochran, Cook, and Dunaway (2000), a starship cover story was used to make the task more interesting and provide meaning to the letter strings. Before beginning the first session, the participants read the following cover story that takes place in a starship:

> We are on a military transport vessel attempting to bring remnants of a space colony back home. Unfortunately, we are short of food for the long trip home. Making matters worse, much of the food that we took on board from the colony has been contaminated by a radioactive poison. Your job is to learn to distinguish poison from non-poisoned food by recognizing poison food labels.
>
> The food taken on board our vessel came originally from another vessel on which all of the passengers died from the poisoned food. Before they all perished, in a last effort to save themselves, members of that ship had installed decontamination devices throughout the ship. These decontamination devices were placed at several control points on the ship where food moved from one location to the next. However, many of the decontamination devices were inoperative. Every can of food that passed through at least one working decontamination device in its travels about the ship, was and still is safe to eat. Cans that passed through only non-working decontamination devices are still poisonous and must not be eaten.
>
> The poisoned food is highly radioactive. Although all of the food supply was initially contaminated, each time it passed through a working decontamination device the amount of radioactivity was reduced. Thus, when tested with a special Geiger counter on the ship, radioactivity levels in individual cans of food may range from 0 to 10. Each can label generated during testing will be located by the computer and tested for radioactivity. Only cans that test at level 10 are poison. Any can with a radioactivity reading lower than 10 is safe to eat. Also, since cans that have readings above 7 are similar to a poison can label (a 10), the computer is capable of tracking down the related poison can and giving the exact label (p. 164-165).

Participants were told that they would see a subset of poison food labels (exemplars) that were saved from destruction. Moreover, they would perform a training task with them that would be useful in discovering more poison food labels during the test. Each participant received a rolodex with a set of exemplars printed on cards and a packet of response sheets for their assigned training task. They were then given a demonstration on how to perform their respective tasks.

Participants in the EP groups were instructed to copy as many of the 88 instances

(exemplars) as possible into the response sheets in 20 minutes. Each letter of each exemplar was to be copied into the appropriate circle on the sheet. Beginning from left to right, participants copied each letter of the exemplar into the circle that intersected the row labeled with that letter and the corresponding column reflecting the ordinal position of that letter within the exemplar (see Panel A Figure 5).

Participants in the ED groups were instructed to trace as many of the 88 poison food labels (exemplars) through the diagrams on their response sheets as possible in 20 minutes. They were instructed to copy each letter of each exemplar into the corresponding transition box until the exemplar was completed (see Panel B of Figure 5). Due to the nature of the grammar, more than one letter can occur at the same transition point. For example, the loops of the grammar allow for certain letters to be repeated, or the switch back toward the end of the grammar that returns to an earlier transition point. When this occurred, participants were instructed to write the letter to the right of the letter(s) already in that box. Participants were shown the proper procedure for tracing an exemplar through the grammar. The exemplar **SCTSSXXVV** was used to demonstrate this task. This exemplar was used because it illustrates the difference between the looping "S" and the recurring "X" and "V". The rationale for this task was to have participants process exemplars within the context of the grammar's structure.

Testing Phase. Participants were told that the ship's computer would display two randomly selected letters and a series of dashes from a not-yet-generated poison food label. Their job was to fill in the dashes with letters that would uncover a poison food label. They worked from left to right in filling in the dashes. When a participant got to a letter that was already revealed, the same letter was retyped. After all the dashes were filled, they pressed the "enter" key. If the letter string generated by the participant did not match at least 70% of the letters of the closest not-yet-generated exemplar, all non-matching letters were erased from the screen and the participant would try again. This process was continued until at least 70% of the letters typed by the participant matched an exemplar. When the 70% criterion was achieved, the computer retrieved the closest not-yet-generated exemplar and displayed it for the participant to observe. Participants then pressed the space bar to begin the next trial with a new test cue.

Because different exemplars may have pairs of letters in common, it was not necessary for the participant to generate the exact exemplar used by the computer to create the two-letter test cue. Thus, participants had some flexibility about which exemplar could be generated on a particular trial. However, once an exemplar was generated, it was removed from the database and could not be generated again during that session. Participants were instructed to find as many poison can labels (exemplars) as possible during the test and encouraged to generate as many perfect exemplars (100% letter match) as possible.

Results

One Way ANOVAs were used to analyze all the data. The results for all four dependent measures are presented in Figure 6. The results on each measure will be discussed in turn.
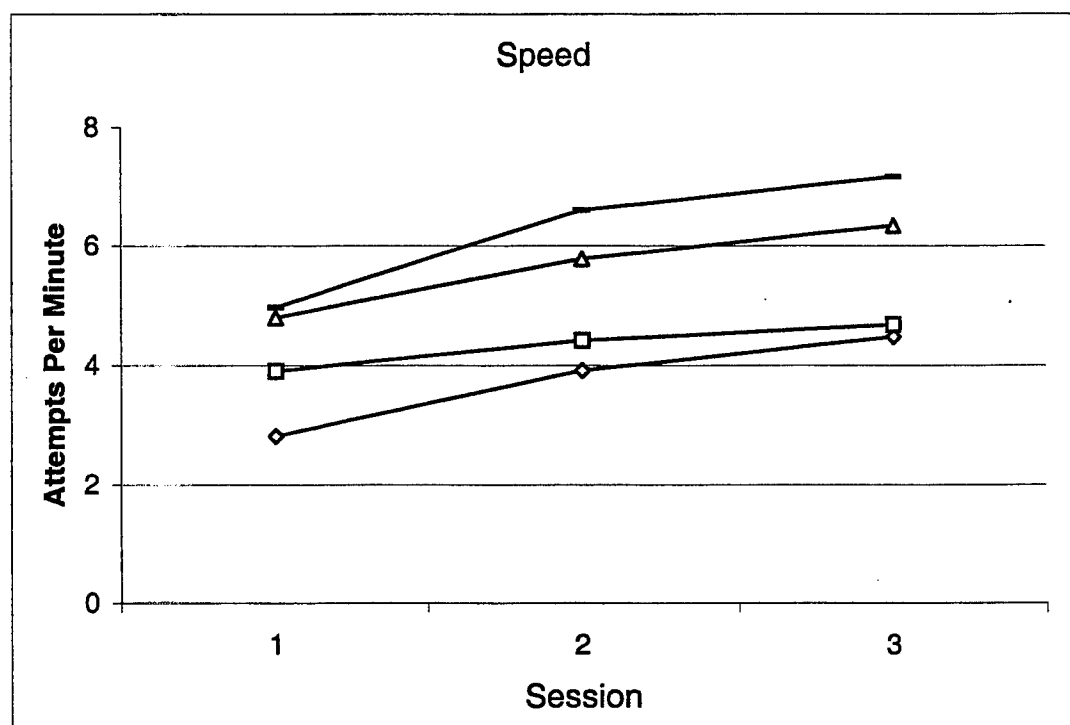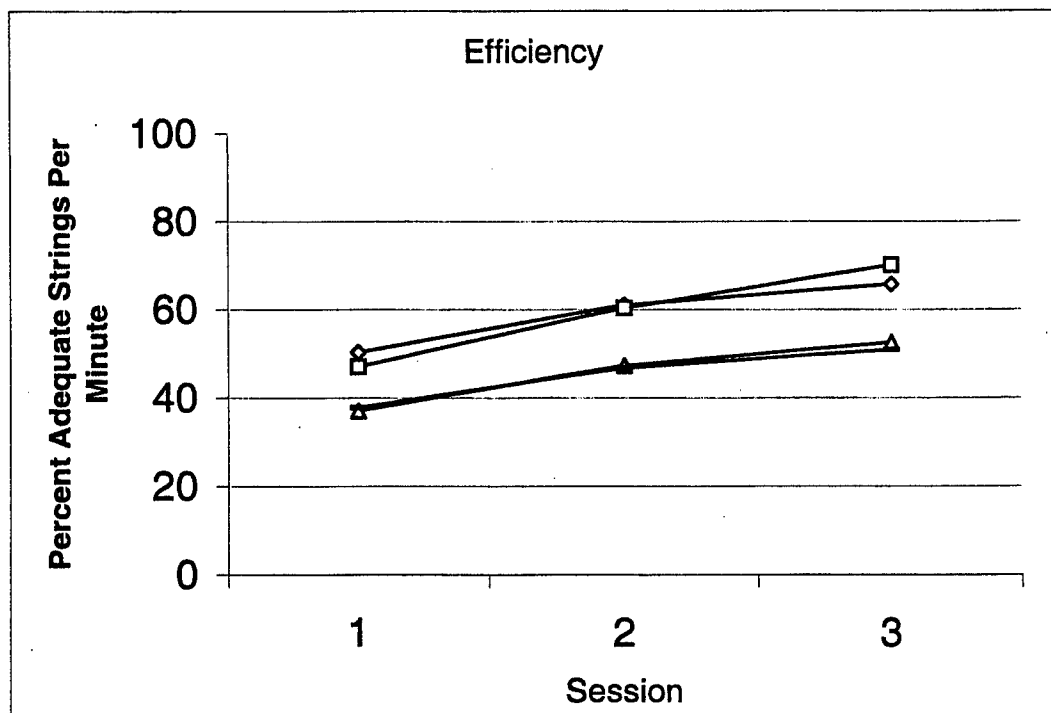
Achievement

Legend: ExD Small, ExD Large, ExP Small, ExP Large

Accuracy
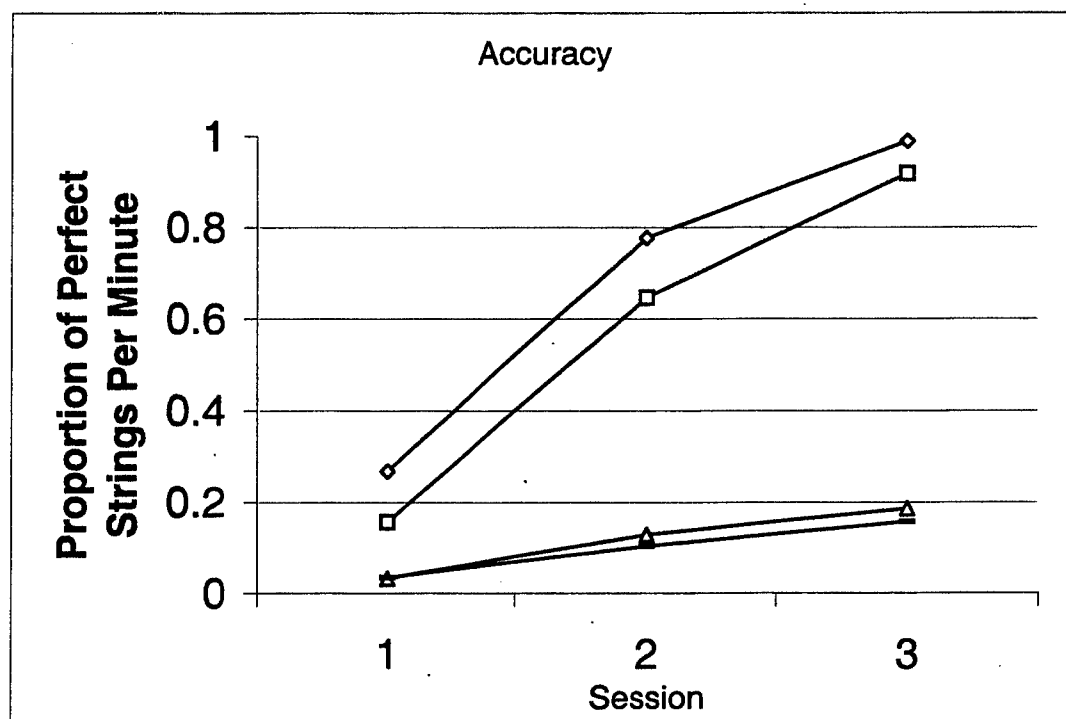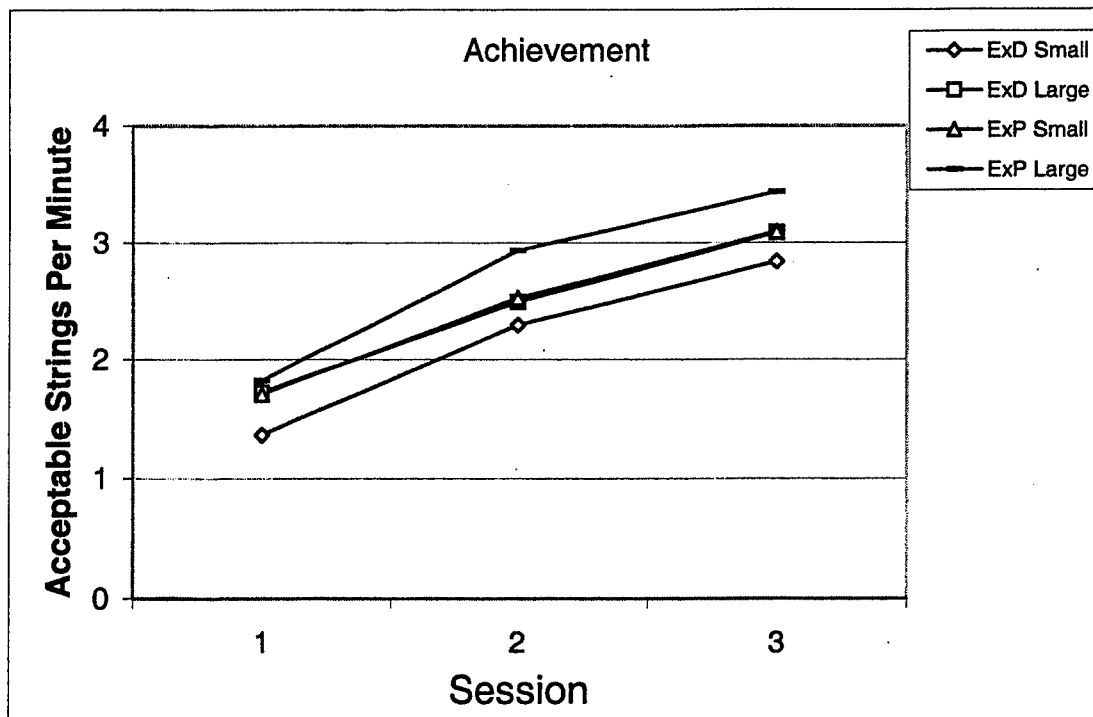
Figure 6. *Results of Artificial Grammar Experiment 1*

Achievement. Achievement is measured in terms of the number of acceptable strings (matching at least 70% of the letters in a not-yet-generated exemplar) successfully generated per minute during the 20-minute test phase. There was a significant effect of sessions on achievement, F (2, 176) = 262.82, MSE = .19, p < .001. Although the achievement levels of all four groups were quite similar (See Figure 6), there was a marginally significant effect of list length, F (1, 88) = 3.41, MSE = 1.58, p = .068, and task, F (1, 88) = 3.56, MSE = 1.58, p = .063. Thus, groups with the large training set achieved slightly more than those with the small training set, and groups with the EP task achieved slightly more than groups having the ED training task. The interaction between list length and task was not significant.

Accuracy. Accuracy is a measure of the proportion of attempts that matched 100% of the letters in a not-yet-generated exemplar (i.e., the proportion of perfect, 100%, letter strings generated per minute). There were significant effects of sessions, F (2, 176) = 27.82, MSE = .17, p < .001 and task, F (1, 88) = 17.30, MSE = 1.07, p < .001. There was also a significant interaction between sessions and task, F (2, 176) = 13.23, MSE = .17, p < .001. Accuracy of the ED groups increased more across sessions than did accuracy of the EP groups (see Figure 6).

Efficiency. Efficiency is a measure of the proportion of a participant's attempts that generate acceptable strings. There were significant effects of session, F (2, 176) = 85.50, MSE = 77.67, p < .001 and task, F (1, 88) = 19.48, MSE = 666.54, p < .001. As can be seen in Figure 6, the ED conditions tended to be more efficient than the EP groups. Also, all groups became more efficient in generating strings across the three sessions.

Speed. Speed of responding was measured in terms of number of attempts per minute during the test phase. An attempt is counted every time the participant pressed the enter key. As expected participants who received explicit training with the grammar (ED task) were slower to respond in the cued-generate test than participants who received implicit (EP task) training. There were significant effects of speed on sessions, F (2, 176) = 79.09, MSE = .72, p < .001, task, F (1, 88) = 27.91, MSE = 8.92, p < .001, and an interaction between sessions and task, F (2, 176) = 3.68, MSE = .72, p = .027. As can be seen in Figure 2, the EP groups performed significantly faster than the ED groups. There was also a three way interaction between sessions, task, and length, F (2, 176) = 5.32, MSE = .72, p = .006. Whereas the EP large group increased in speed over sessions more than the EP small group, the opposite pattern was observed for the ED groups (see Figure 6).

Discussion

The results of the first experiment of this series demonstrate that there are both advantages and disadvantages of exposing participants to an explicit representation of the grammar during training. Explicit knowledge of the grammar acquired in the ED groups led to better accuracy in terms of generating more perfect strings. It also led to greater efficiency in terms of the proportion of strings generated that were acceptable in the cued-generate test (matching at least 70% of the letters in a not-yet-generated exemplar). However, the EP groups, who did not have this explicit knowledge, responded faster, allowing them to generate more valid strings during the 20 minute test. These results support the view that purely implicit knowledge acquired from processing exemplar strings is sufficient to support generation of acceptable (70% correct) strings.

There was also a marginal effect of training set size on achievement (number of strings generated). Groups that received the large training set (88 different exemplars) generated slightly more strings than groups that received the small training set (22 exemplars randomly repeated four times). However, this effect was very small. Thus, an extensive memory bank of exemplars does not appear to be necessary for learning an artificial grammar.

Introduction

In some past experiments researchers have found that mixing different training tasks across sessions could enhance learning. Experiment 2 examines inter-session mixing of the two training tasks (EP and ED).

In Experiment 1 we found that implicit training (EP task) led to the fastest responding on the cued-generate test. However, explicit training, processing exemplars within the context of the grammar diagram (ED task), led to greater accuracy and efficiency in generating strings. A few previous experiments (Reber et. el., 1980; Mathews et. al., 1989) have examined mixtures of implicit and explicit training across sessions, and found mixtures to be more effective than receiving a single training task. However, these studies differed in terms of which combination was best, and neither of the studies examined performance in a task that involves both speed and accuracy.

This experiment examined the effects of mixing EP training with ED training across two weekly sessions. Perhaps groups with mixed training (EP,ED or ED,EP) would acquire the best qualities of both types of training, faster than ED and more accurate than EP. Experiment 2 also included a one-week retention test without a training phase during the third session. This retention test was included because it has often been found that conditions which lead to the fastest initial learning do not usually result in the best retention (e.g., Pollock & Lee, 1997; Shewokis, Del Rey, & Simpson, 1998). It was predicted that the group receiving ED training during the first two weekly sessions would perform best in retention since these participants should have retained a visual representation of the grammar in addition to their implicit memory bank of instances.

Method

Participants. One hundred eight undergraduate students taking a variety of psychology courses at Louisiana State University participated in the experiment. All participants were volunteers and received extra credit for their participation. None of the participants from Experiment 1 participated.

Materials. The same materials from Experiment 1 were used in this experiment with the exception of the elimination of the large set of training exemplars.

Design. The design was a one-factor between-subjects design with four levels: EP during the first two sessions, ED during the first two sessions, EP during the first session and ED during the second session, and ED during the first session and EP during the second session. Twenty-seven participants were randomly assigned to each of the four conditions.

Procedure. The procedure was exactly like the first experiment in all aspects except two. The first was that two groups received a different training task during their second session than they did during the first session (i.e., mixed groups). The second was that participants did not perform any training task during their third session. Instead, during the third session, they performed the cued-generate test for 40 minutes. The test time was increased in the retention

session to obtain a more thorough assessment of participants' ability to generate a wide range of valid strings after the one-week retention interval.

Results

The data from all three sessions (including retention) are shown in Figure 7.

Accuracy

Efficiency
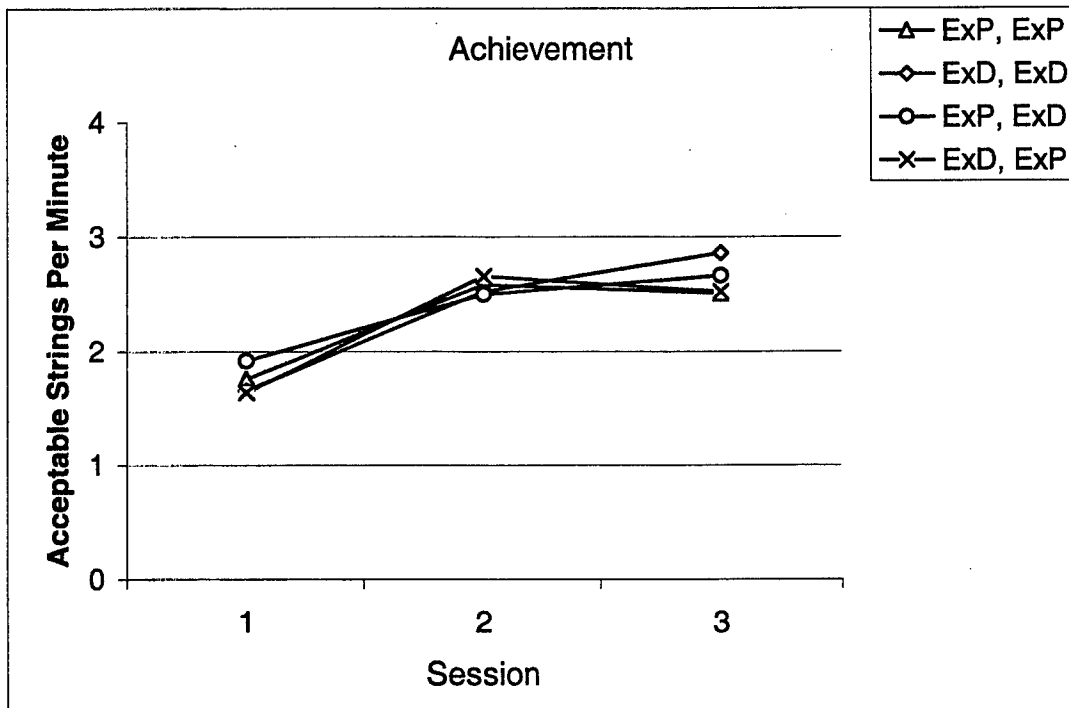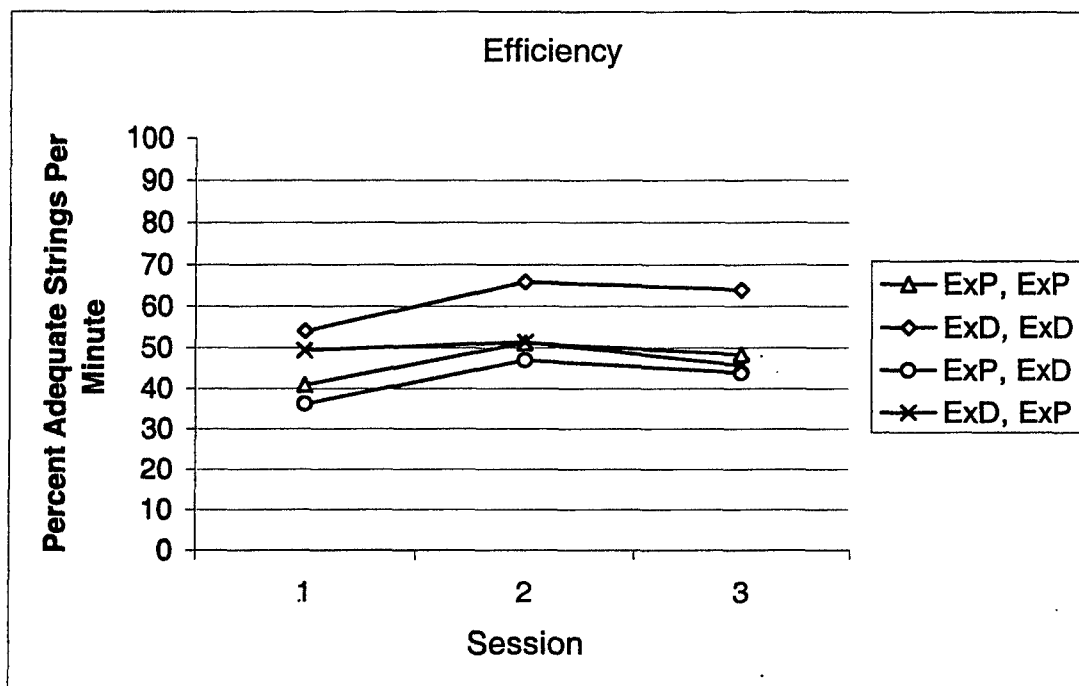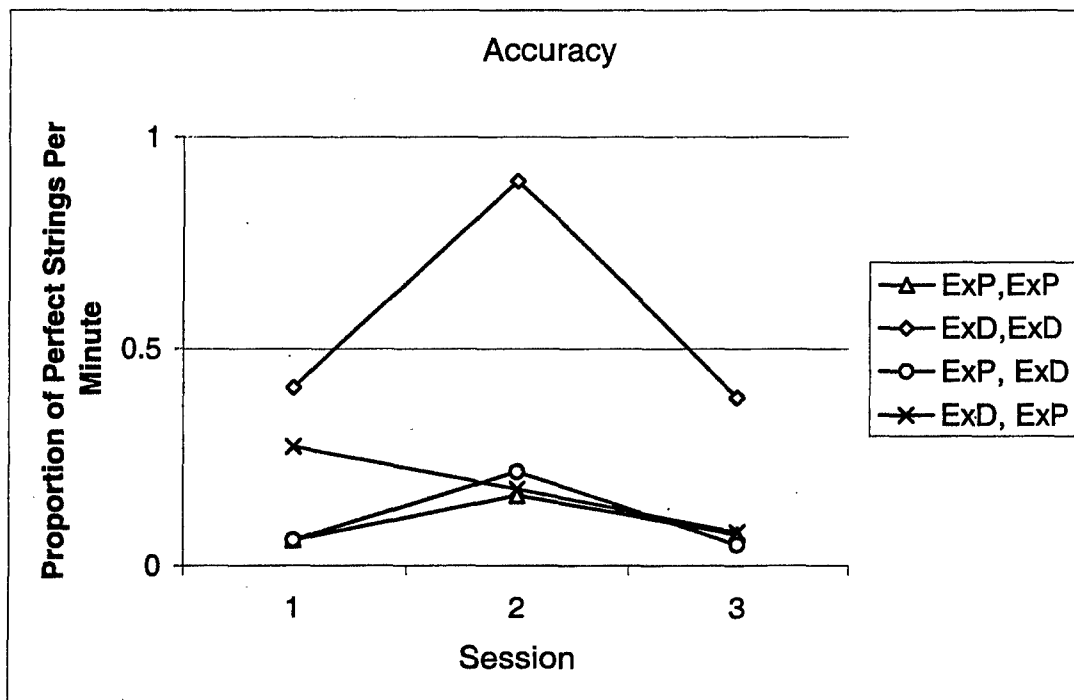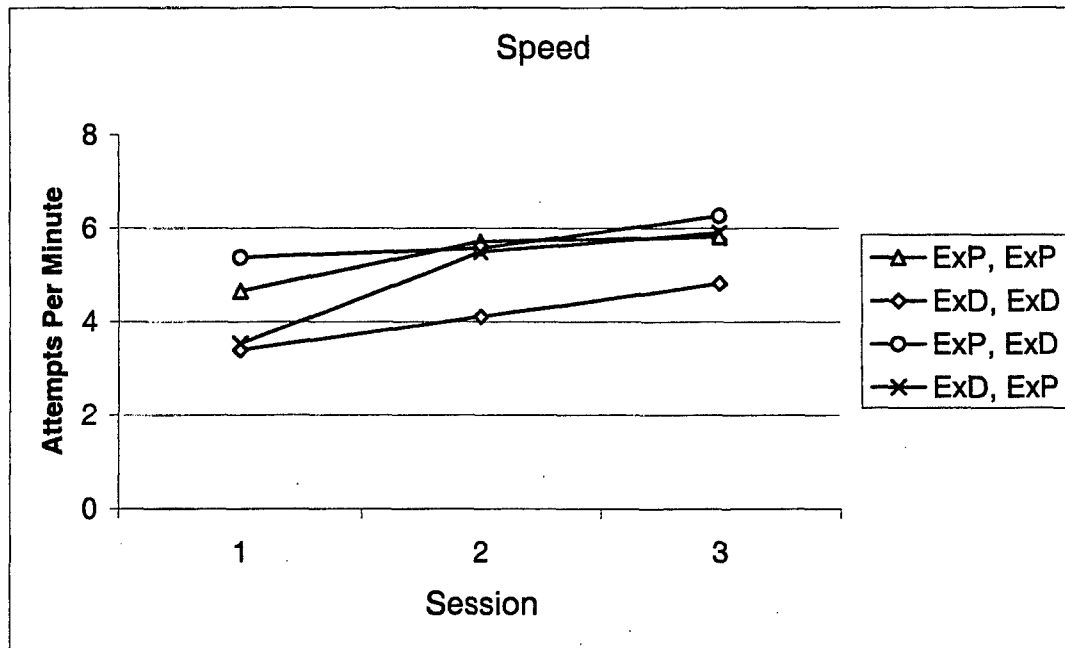
Figure 7. *Illustrates the performance during Experiment 2 of various training tasks on the four dependent measures. Also, the exemplar processing (EP) training and the exemplar diagramming (ED) training were mixed from week one to week two (EP followed by ED and ED followed by EP). The third session contained no training phase and extended the cued-generate test from 20 minutes to 40 minutes (testing over a retention interval).*

The data from the second session and the retention session are of primary interest because the mixed groups have not experienced both types of training until the end of session 2. Also, recall that the test phase during the retention session was twice as long (40 minutes) as the test during acquisition (20 minutes). This additional time was provided to determine if performance levels could be maintained when participants were required to generate a greater number of valid strings. Doubling the length of the retention test (40 minutes instead of 20) would permit participants to generate twice as many strings if they maintained their levels of speed and accuracy during the extra 20 minutes of the retention test. Given the different amount of time allowed for the test, the acquisition data (Session 2) and retention data (Session 3) were analyzed separately and will be discussed in turn.

Acquisition Phase Analyses

Achievement. There was no significant effect of training tasks on achievement during acquisition.

Accuracy. There was a significant effect of training tasks on accuracy during acquisition, $F(3, 104) = 7.65$, $MSE = .45$, $p < .001$. A *Tukey HSD* post hoc test of comparisons showed that the ED, ED group ($M = .90$) was significantly more accurate than all other groups, which did not differ from each other.

41

Efficiency. There was a significant effect of training task on efficiency during acquisition, F (3, 104) = 5.67, MSE = 326.88, $p < .001$. A *Tukey HSD* post hoc test of comparisons showed that the ED, ED group (M =65.88) was significantly more efficient than all other groups, which did not differ significantly from each other.

Speed. There was a significant effect of training task on speed during acquisition, F (3, 104) = 3.69, MSE = 4.06, $p = .014$. A *Tukey HSD* post hoc tests of comparisons showed that the EP, EP group ($M = 5.71$) and the EP, ED group ($M = 5.57$) performed significantly faster than the ED, ED group ($M = 4.11$). The ED, EP group did not differ significantly from any other group.

Retention Phase Analyses

Achievement. There was no significant effect of the training tasks on achievement during retention. All groups were able to maintain their level of achievement on the extended (40 minute) retention test. Note that all groups generated approximately the same number of strings per minute in the longer retention session as compared to the 20 minute acquisition test (see figure 3). Thus, the rate of generating valid strings did not diminish in the extended retention test.

Accuracy. There was a significant effect of training task on accuracy during retention, F (3, 104) = 9.73, MSE = .07, $p < .001$. A *Tukey HSD* post hoc test of comparisons showed that the ED, ED group ($M = .39$) was significantly more accurate at string generation after a one-week retention period than all other groups, which did not differ from each other. However, it should be noted that the ED, ED group showed the largest drop in accuracy from Session 2 to Session 3 (See Figure 3). This result was surprising because we expected that having both implicit and explicit knowledge of the grammar would enhance retention.

Efficiency. There was a significant effect of training task on efficiency during retention, F (3, 104) = 9.73, MSE = 234.56, $p < .001$. A *Tukey HSD* post hoc test of comparisons showed that the ED, ED group ($M = 64.01$) was significantly more efficient after a one-week retention period than all other groups, which did not differ from each other.

Speed. There was an effect approaching significance of training task during retention, F (3, 104) = 2.31, MSE = 4.48, $p = .08$. The ED, ED group performed slower than all other groups.

Discussion

As in Experiment 1, all types of training led to similar levels of achievement on the cued-generate test during both acquisition and retention. Interestingly, the mixed groups performed more like the implicitly (EP) trained groups, responding quicker, but with less accuracy and efficiency as compared to the ED groups. This pattern of results suggests that exposure to implicit training either before or after explicit training led our participants to prefer their implicit (fast but less accurate) mode of responding to the task. Perhaps this is because using the explicit knowledge of the grammar is effortful and slow. Participants seem to be naturally drawn to the implicit mode in this task because perfect accuracy was not required (computer motherese was available). Moreover, these patterns were maintained during the one-week retention interval.

In a sense the ED training task is a mixed (implicit and explicit) form of training. Participants having this training task process exemplars (implicit training) in the context of a

42

diagram of the grammar (explicit training). The final experiment of this series adds another training task that is closer to being purely explicit. This new type of training task, called grammar reproduction or GR, requires participants to commit to memory the diagram of the grammar without processing exemplars during training. Experiment 3 also examined mixes of this new more explicit (GR) training task with the purely implicit (EP) training task.

Artificial Grammar Experiment 3

Introduction

In Experiment 3, the EP training task continued to serve as the implicit training task, while a new training task was created to provide explicit training without opportunities to process many exemplars (controlling implicit contamination). The new task was termed "grammar reproduction" (GR). Very few experiments have provided participants with the grammar diagram during training. In the few studies that have provided such explicit knowledge of the grammar, it was provided for a very minimal amount of time (e.g., Reber et. al., 1980). In this experiment GR trained participants committed the entire diagram to memory before attempting to generate strings.

It was predicted that participants having only the purely implicit (EP) training would generate strings the fastest, using only fast implicit processes. It was expected that the purely explicitly (GR) trained group would be the most accurate, but the slowest, using only explicit knowledge. The integrated (ED) training was expected to fall in between the two pure groups, employing some fast implicit processes combined with slower explicit knowledge. We also examined mixed GR and EP training across sessions to see which type of training produced optimal results for combining implicit and explicit processes. A control group was also added to explore performance in the absence of any type of training task. Although this group had no training, they were expected to perform above chance on the cued-generate test because they could rapidly type each of the six possible letters in succession until a 70% match was obtained. Miller (1969) termed this a cyclic strategy. Thus, the control group might do well in achievement, but their efficiency and accuracy measures were expected to be very low.

Method

Participants. One hundred twenty undergraduate students taking a variety of psychology courses at Louisiana State University participated in the experiment. All participants were volunteers and received extra credit for their participation. No participant from the two previous experiments participated in Experiment 3.

Materials. The same materials used in Experiment 2 were used in this experiment.

Design. The design was a one-factor between-subjects design with six levels: EP during both weeks, ED during both weeks, grammar reproduction (GR) during both weeks, EP followed by GR, GR followed by EP, and a no training control (C) during both weeks. Twenty participants were randomly assigned to each of the six conditions.

Procedure. There were two 1-hour sessions conducted one week apart with a 20 minute training phase and a 20 minute testing phase. Participants followed the same instructions from the prior experiment for performing the EP and ED tasks. The GR training task required participants to observe a copy of the artificial grammar for 2 ½ minutes then turn the diagram over. For another 2 ½ minutes participants reproduced the artificial grammar diagram from memory by drawing it on a blank sheet of paper. This was repeated four times for a total of 20 minutes training time, consistent with the other training tasks.
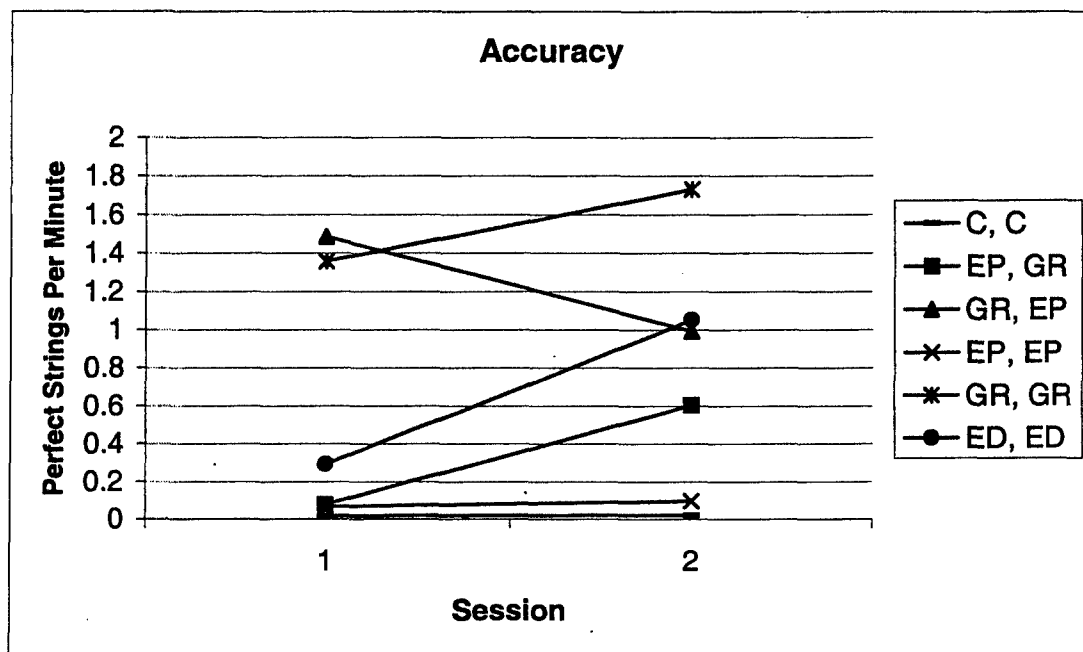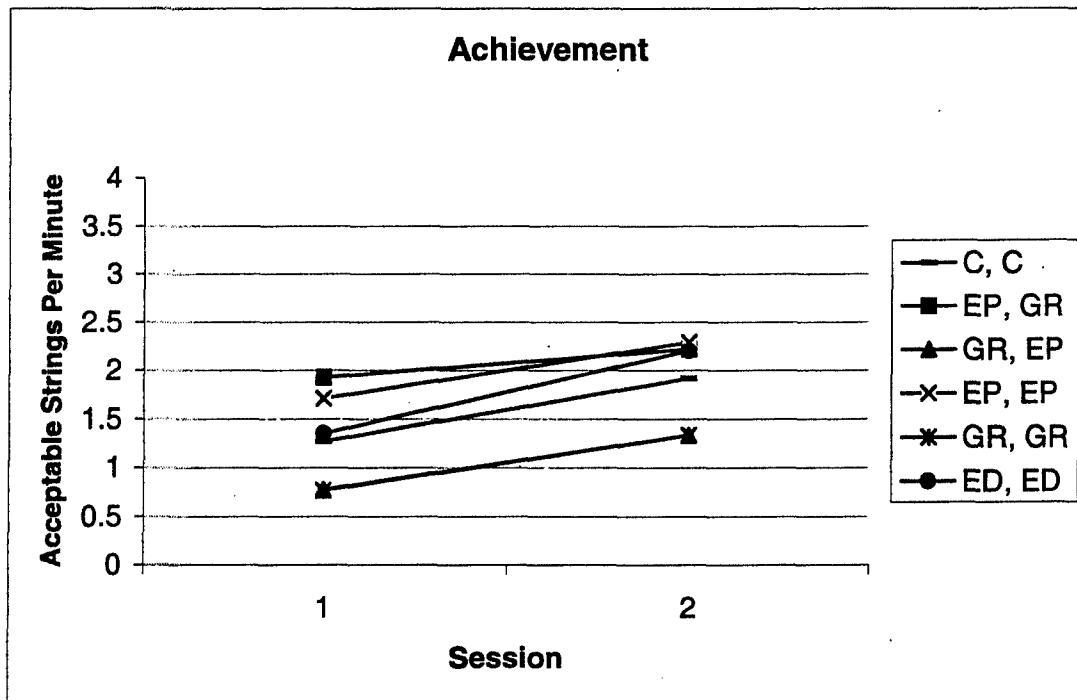
44

The goal of the GR task was to teach an explicit representation of the grammar without showing many valid letter strings that could stimulate implicit learning. However, it was essential that participants understood how to use the diagram to generate strings. Therefore, prior to the first session, three test cues were used to demonstrate how to generate strings using the diagram. One test cue used for this purpose was - - T X -. Participants were shown how to generate two different valid strings, **SCTXS** and **CXTXS**, using this cue. The second string demonstrated was - - P - - P -. The strings **SCPTVPS** and **CXPTVPS** were generated from these cues. The third string demonstrated was - - - T - - - X - -. In this case, only one the exemplar, **CVCTSSXXVV** can be generated. These cues, increasing in complexity, demonstrated some of the properties of the grammar such as the fact that a letter can occur twice (e.g. both the "X" and the "V") without being in a loop.

The control (C) condition did not receive any training. They were given a sheet of paper with the six letters of the grammar, typed in 36 point Courier font, randomly placed horizontally across the middle of the page. The only instructions given to these participants were to try and generate letter strings by filling in the blanks by typing combinations of the six letters of the grammar and press enter. Correct letters would remain on the screen and should be used in combination with other choices for another attempt until an acceptable string is generated. They were also informed about the 70% minimum criterion and the ability of the computer to provide the corrected exemplar.

The 20 minute testing phase was identical to the prior experiments.

Results

Only the results from the second session were analyzed statistically because the mixed groups (EP,GR and GR,EP) did not experience both training tasks until the end of the second session. However, performance measures for both sessions are provided in Figure 8. Figure 8. Illustrates the performance during Experiment 3 of various training tasks on the four dependent measures. The grammar replication (GR) training task was implemented and also mixed with the EP training task. A control (C) condition, which received no training, was also added.

Achievement

Acceptable Strings Per Minute vs. Session

Legend: C, C | EP, GR | GR, EP | EP, EP | GR, GR | ED, ED

Accuracy

Perfect Strings Per Minute vs. Session

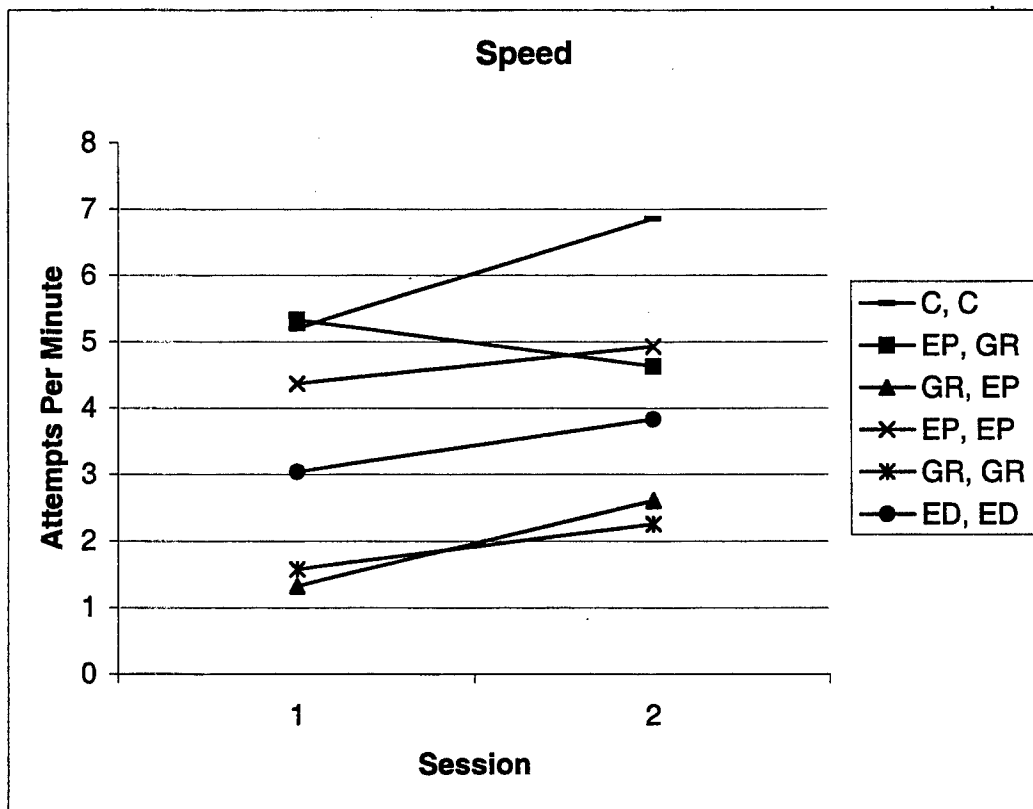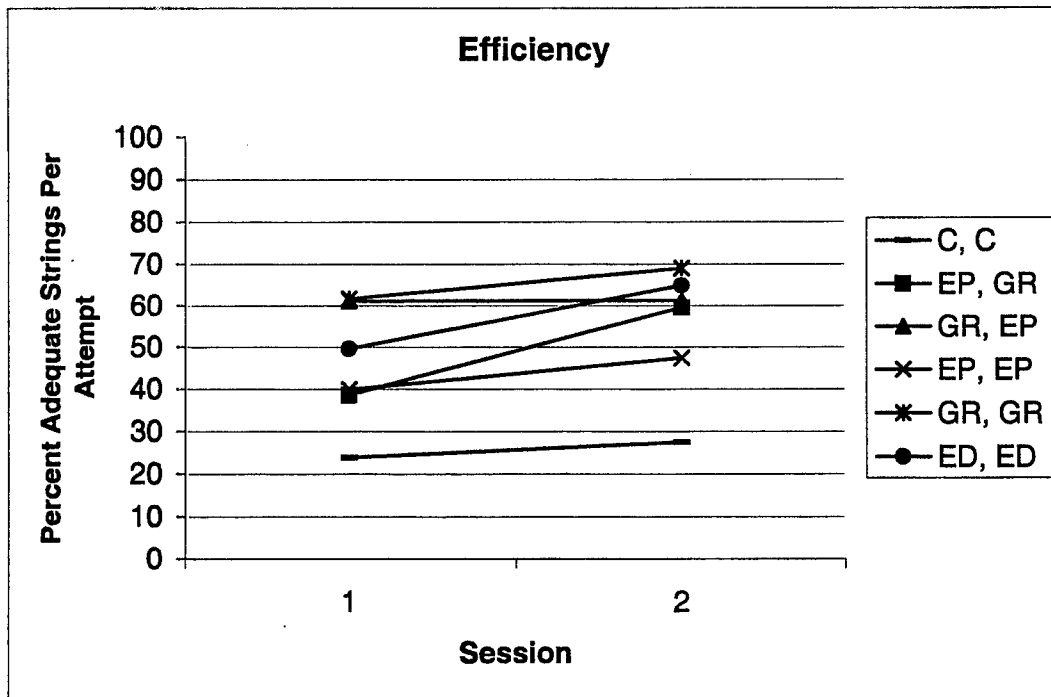Legend: C, C | EP, GR | GR, EP | EP, EP | GR, GR | ED, ED

Figure 8. *Results of Artificial Grammar Experiment 3*

Achievement. There was a significant effect of training tasks on achievement, F (5, 114) = 6.81, MSE = .58, p < .001. A *Tukey HSD* post hoc test of comparisons showed that the GR, EP group (M = 1.34) and the GR, GR group (M = 1.34) performed significantly less well than all other groups except for the C, C group (M = 1.92) which did not differ significantly from any other group.

Accuracy. There was a significant effect of training tasks on accuracy, F (5, 114) = 8.82, MSE = .95, p < .001. A *Tukey HSD* post hoc test of comparisons showed that the GR,GR group (M = 1.73) was significantly more accurate than the C, C group (M = .02), the EP,EP group (M = .10), and the EP, GR group (M = .60) which did not differ significantly from each other. The EP, GR group only differed significantly from the GR, GR group.

Efficiency. There was a significant effect of training tasks on efficiency, F (5, 114) = 10.03, MSE = 466.50, p < .001. A *Tukey HSD* post hoc test of comparisons showed that the C, C group (M = 27.54) performed significantly worse than all other groups while the GR, GR group (M = 69.04) was significantly more efficient than the C, C group and the EP, EP group (M = 47.41). The EP, EP group only differed significantly from the C, C group and the GR, GR group.

Speed. There was a significant effect of training tasks on speed, F (5, 114) = 14.11, MSE = 4.02, p < .001. A *Tukey HSD* post hoc test of comparisons showed that the C, C group (M = 6.86) was significantly faster than all other groups. The GR, GR group (M = 2.26) and the GR, EP group (M = 2.61) were significantly slower than all other groups except for the ED, ED group (M = 3.83) which only differed from the C, C group.

Discussion

Experiment 3 compared purely explicit training (GR) to purely implicit training (EP) and integrated training (ED). It also examined various mixtures of training type across two sessions. The results followed the pattern of the earlier experiments in that exposing people to a diagram of the grammar (GR or ED) generally led to slower but more accurate responding on the cued-generate test. Memorizing the grammar without encoding exemplars during training (GR) led to the highest level of accuracy and the slowest responding. Purely implicit training led to fast responding with low accuracy. The integrated training was in between, having higher accuracy and lower speed than EP, and lower accuracy but higher speed compared to GR.

Whereas, in the earlier experiments, achievement (number of strings generated) was nearly equivalent across groups, in this experiment large differences occurred. The pure explicit (GR, GR) group had lower achievement, even compared to the control group who had no training. However, the implicitly trained (EP, EP) group and the integrated training (ED, ED) group were able to generate more strings than the explicitly trained (GR, GR) group or the control (C, C) group in Session 2.

Interestingly, the groups exposed to mixed training across sessions tended to perform like the pure groups who had similar training in Session 1. Consequently, the GR, EP group did poorly on the achievement measure (as did GR, GR); and the EP, GR group successfully generated as many strings as the EP, EP group. Thus, it appears that the type of training received initially tends to dominate when training type is changed.

48

In this section we simulated our human data from Experiment 3 with CLARION, an integrative model with a dual representational structure (Sun et al., 2001; Sun, 2002). As mentioned before, the model consists of two levels: the top level encodes explicit knowledge and the bottom level encodes implicit knowledge. The purpose of the simulation was to see if a model using dual representational structures could capture the key features of our data. No attempt was made to fine tune the fit of the model by varying parameters, because at this stage we are only interested in the overall features of the data.

As mentioned before, the inaccessible nature of implicit knowledge is captured by the subsymbolic distributed representations provided by a backpropagation network (Rumelhart et al., 1986). This is because representational units in a distributed representation are capable of accomplishing tasks but are subsymbolic and generally not individually meaningful (see Rumelhart et al., 1986; Sun, 1994); that is, they generally do not have an associated semantic label. This characteristic of distributed representation accords well with the inaccessibility of implicit knowledge.

In contrast, explicit knowledge may be captured in computational modeling by a symbolic or localist representation (Clark & Karmiloff-Smith, 1993), in which each unit is easily interpretable and has a clear conceptual meaning (i.e., a semantic label). This characteristic captures the property of explicit knowledge being accessible and manipulable (Smolensky, 1988; Sun, 1994).

This radical difference in the representations of the two types of knowledge leads to a two-level model whereby each level using one kind of representation captures one corresponding type of process (either implicit or explicit). The model may select to use one level or the other, based on current circumstances (e.g., experimental conditions; see Sun, 2002 for details). When both levels are used, the outcome from the two levels may be combined through some stochastic selective processes that may be partially domain specific (Sun, 2002).

At each level of the model, there may be multiple modules, both action-centered modules and non-action-centered modules (Schacter, 1990; Moscovitch & Umilta, 1991). The reason for having both action-centered and non-action-centered modules at each level is because action-centered knowledge (roughly, procedural knowledge) is not necessarily inaccessible directly, and non-action-centered knowledge (roughly, declarative knowledge) is not necessarily accessible directly. Although it was argued by some that all procedural knowledge is inaccessible directly and all declarative knowledge is directly accessible, such a clean mapping of the two dichotomies is untenable in our view.

At the bottom level of the non-action-centered subsystem, experienced strings (as presented to subjects or sampled from presented grammar diagrams) are used to train an associative memory made up of a backpropagation network. The network maps input to output; in this particular case, it maps some partial strings (each of which is a part of an experienced string) to the full experienced string. This associative mapping allows implicit grammatical knowledge to develop. This method of training can be justified based on the fact that such associative learning can be easily performed from observing a given string and it can provide the needed implicit grammatical knowledge (as embedded in the network weights).

49

At the top level, experienced strings are encoded as associative rules. For example, if a string ``S C P V" is experienced, the following three rules may be encoded there: S->C, C->P, P->V.

The outcome from the model can be either from the bottom level or from the top level. However, the bottom-level implicit processes are significantly faster than the top-level explicit processes (see Schneider & Oliver, 1991; Hunt & Lansman, 1986; Sun & Zhang, 2001). In CLARION, response time is determined by parameters that specify the time lag of each step of associative memory retrieval at the bottom level, and the time lag of each step of rule application at the top level.

For the explicit/explicit (GR,GR) group, the top level is mainly responsible for generating the outcome during test. This is because, given the initial experimental setting during training, the system was configured in such a way that mainly the top level is used, due to the fact that this experimental setting encourages an explicit mode because of the presentation of grammar diagrams (and thus grammatical structures) to subjects during training. The cross-level combination parameters were automatically set during training in a way that supports this configuration. During test, the top level uses learned rules to attempt to complete each given partial string. That is, given the test cue, it searches for a possible completion guided by the rules at the top level, using depth-first search with backtracking. For example, given a partial string ``S _ _ V", the search has to go through all the rules in the form of S->x, or in the form of x->V, where x can be any letter, and many other similar rules (e.g., concerning the relation between the second and third letters). This search process is slow, but the outcome from the top level is rather accurate. When a completion of a partial string is found, and it is completely consistent with the rules available, the completed string is used as output. However, if a completion is impossible using given rules at the top level (due to the lack of applicable rules), the model attempts to complete as many positions as possible (it compares different partial completions and chooses the most complete one). Then, the bottom level is used. The partially completed string generated thus far by the top level is used as input to the bottom level to come up with a full string. Then, this (guessed) completion is used as output.

For the implicit/implicit (EP, EP) group, during test, the bottom level is responsible for generating the outcome. This is because, given the experimental setting during training, the system is configured in such a way that mainly the bottom level is used, due to the fact that this experimental setting during training encourages an implicit mode, through repeatedly presenting training instances. The cross-level combination parameters were automatically set during training in a way that supports this configuration. During training, the bottom level uses an associative memory (in the form of a backpropagation network) to map a given partial string (test cue) to a full string that is a likely completion of the partial string. This way of capturing implicit learning during training is especially appropriate, considering the fact that subjects in this task marked experienced strings on a bubble sheet, which naturally led to multiple partial strings. The bottom level is, generally speaking, less accurate but much faster.

For the integrated training (ED,ED) group, a combination of the two levels was used, because the experimental settings involve both implicit training and explicit training, due to the use of both repeated presentation of strings and the presentation (and tracing) of grammar diagram. During test, the combination process of the two levels proceeds this way: The bottom level generates candidate completions of partial test strings; then the top level checks each of

these strings using the rules already learned at the top level. The check by top-level rules is carried out through straightforward application of relevant rules, without any backtracking. For example, if ``S C P V'' was suggested by the bottom level, at most three rules may be applied: S->C, C->P, P->V, if these rules do exist at the top level. Thus, in this case, the top level works faster than that of the explicit/explicit (GR,GR) group because in the latter case, there is no suggested string from the bottom level that is available.

If all the relevant rules are available and consistent with the candidate completion of the given partial string as generated by the bottom level, then that completion is used as output. If any of these rules are absent, an alternative rule will be used, which corrects the position that failed validation. In this case, although the bottom level works at a fast pace, the top level is slower. But because there is no full-blown depth-first search with backtracking, the top level is not as slow as in the case of the explicit/explicit (GR,GR) group. But due to multiple applications of rules, it is definitely slower than the bottom-level implicit processes alone. So, the final outcome is, on average, at a speed somewhere between the implicit/implicit (EP,EP) group and the explicit/explicit group.

We made the simplifying assumption that the implicit/explicit (EP,GR) group is essentially the same as the implicit/implicit group in terms of using mainly the bottom level in generating responses during test. This is because the initial implicit experimental setting during the first training session may have locked that group into using mainly the bottom level the same way as the implicit/implicit group. The cross-level combination parameters were set during the first training session, which are unlikely to change.

Likewise, we made the simplifying assumption that the explicit/implicit (GR,EP) group is essentially the same as the explicit/explicit group in terms of using mainly the top level in generating responses. This is because the initial explicit experimental setting during the first training session may have locked that group into using mainly the top level in ways similar to the explicit/explicit group.

To model the control/control (C,C) group, no training was done. The bottom level is used to generate responses. The associative memory produces essentially random guesses (due to the lack of training).

The training of the bottom level, the encoding of rules at the top level, and the selection of outcomes from either level, the search at the top level to generate a completion or to validate a candidate completion are all under the control of the actions by the action-centered subsystem (ACS). It makes action decisions each step of the way, in sequential order. Thus, the ACS directs the operation of the non-action-centered subsystem. Details regarding the ACS and its parameters, and the details of how it directs the NACS, are omitted here due to their complexity (see Sun, et al., 2001; and Sun, 2002 for more detailed descriptions). The dependent variables are essentially parallel to those obtained from the human data.

The key features we were trying to capture in the simulation were that exposure to a diagram of the grammar either through grammar replication (GR) or exemplar diagramming (ED) would enhance accuracy and efficiency but such exposure would reduce speed. Plus, a high level of achievement could be accomplished through implicit processing (EP,EP) alone, without exposure to a diagram of the grammar. The results are shown in Figure 9.

Note that simulation outcomes of different groups vary because of a number of independent factors: cross-level combination differences in generating responses (e.g., relying on

the bottom level vs. relying on the top level in generating responses), training differences (e.g., due to different training data used in EP vs. GR), random variations (e.g., due to random initializations of weights in backpropagation networks). The results in Figure 9 should be viewed in this light.
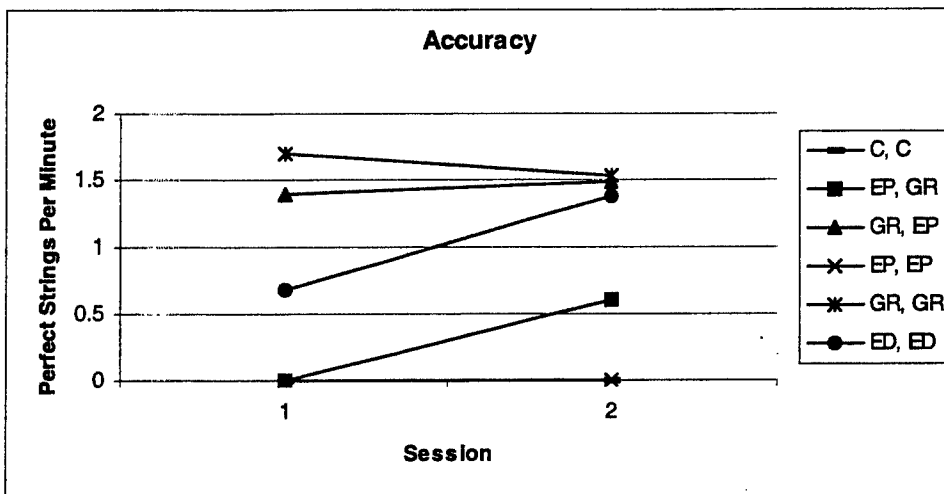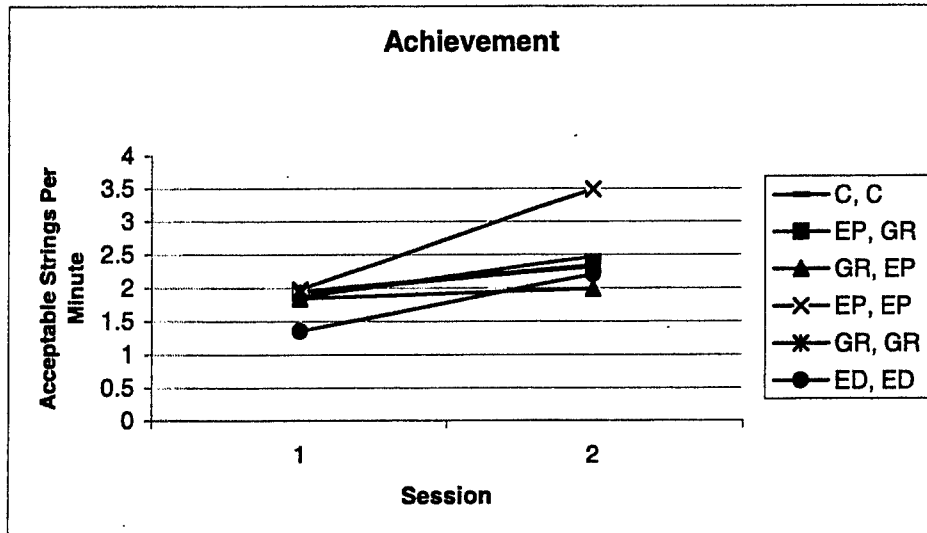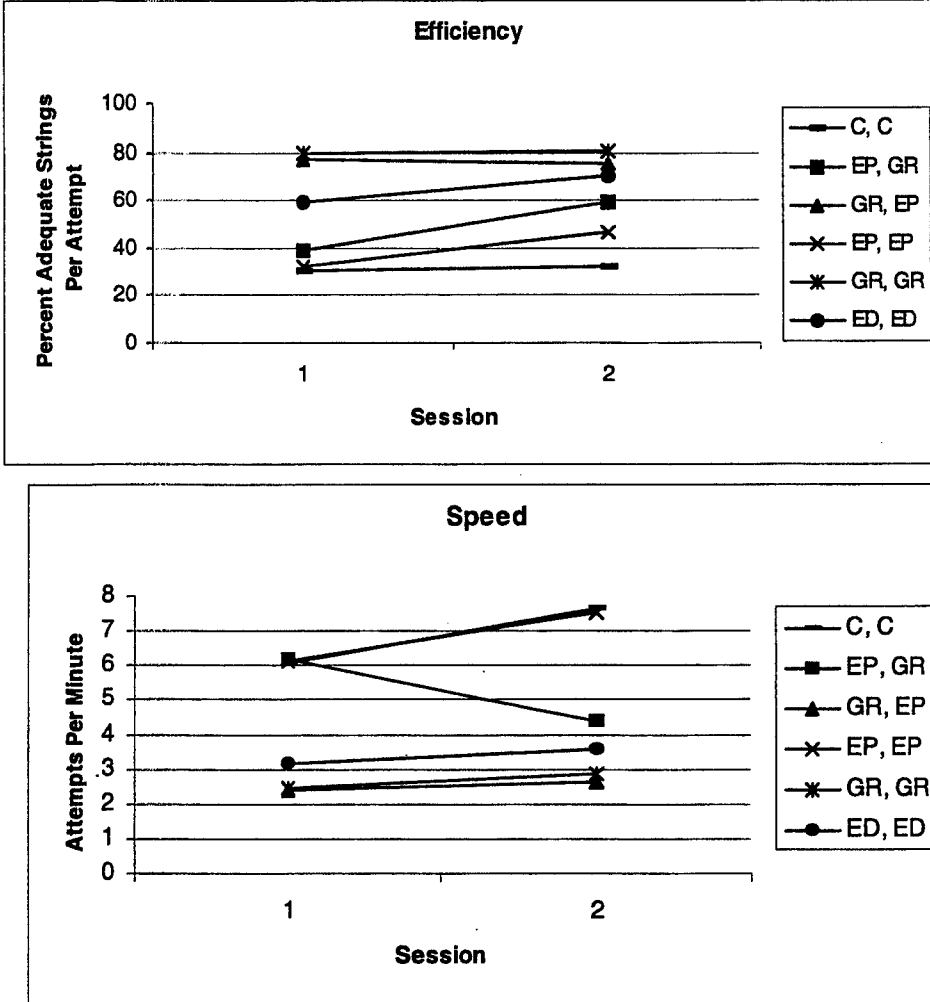
Figure 9. *Experiment 3 simulation illustrates the performance of the CLARION model capturing the human data from Experiment 3.*

The simulation results for the accuracy and efficiency data are quite similar to the human data. In Week 2, the three highest groups in both the simulation and human data were those exposed to a diagram of the grammar. However, in the human data, the grammar replication (GR,GR) group was superior to all other groups, while in the simulation it was only slightly better. As expected, the groups exposed to the diagram were more efficient in both the human and simulation data.

However, exposure to the diagram also reduced the speed of string generation. In both the human and simulation data, the control group (C,C) and the group only exposed to implicit training (EP,EP) were fast. However, in the human data, but not in the simulation, the control group was faster than the implicitly trained group. Finally, the implicit only group (EP,EP) had a high level of achievement in both the simulation and the human data.

Thus, the simulation results support the notion that a dual representational model can account for these data. Further studies and simulations are planned using reaction time data to study different ways knowledge in the two levels can be strategically applied to a task.

Our simulation using CLARION has produced some interesting interpretations of the human data. These interpretations are embodied in our simulation setups as described earlier. They described a plausible mechanistic underpinning of human performance in this task. In particular, they provide an explanation of why the integrated training group performed better in Experiment 3.

Introduction

The integrated training conditions in the preceding experiments used a form of practice in which learners mapped exemplars into a diagram of the grammar. This task can be characterized as parsing whole exemplars into parts and placing them within the structure of the grammar diagram. While this type of training provides insight into how exemplars are constructed, it might reduce attention to whole exemplars. If storage of intact whole exemplars is important to the implicit learning process (Brooks, 1978; Whittlesea & Wright, 1997), this might not be the optimal form of training for integrating implicit and explicit learning. The remaining two experiments changed the practice task so that the emphasis is on processing whole, intact exemplars. However in the integrated training condition, an animated form of the explicit grammar diagram is used to prime encoding of the exemplar.

Experiment 4 compared performance in a transfer task involving string generation following training. Training was conducted through the use of three different computer games in which participants performed a string edit task. The goal of all three training games was the same: participants were shown a letter string and told to identify the incorrect letters in that string. Their "score" was presented in terms of misses (incorrect letters that they did not identify as such) and false alarms (correct letters identified as incorrect). Participants were encouraged to make few errors and a monetary prize was offered to the participant in each condition who made the fewest errors. While the goal of the games was the same, they differed in the type of assistance given to the participant.

Participants in the letter appearance (LA) condition attempted to identify the incorrect letters in the string without any assistance. They were shown a letter string at the bottom of the computer screen and told to select the incorrect letters in that string and click on them with the mouse. As the trial progresses, the computer presents the correct string at the top of the screen, with each letter appearing one-by-one from left to right, until the entire string is revealed. Approximately 3 seconds after the trial begins, letters begin appearing at the top of the screen, and 500 ms before a letter appears in its position at the top of the screen, participants can no longer edit the letter in that position. Thus, participants are required to make fairly quick decisions.

In the primed assist (PA) condition, participants were given the same string-edit task as the LA condition, but were provided an aid to prime correct choices. Instead of the correct letters appearing one-by-one at the end of the trial as in the LA condition, the letters emerged from an unrecognizable bunch in the bottom of the screen and became recognizable as they slowly floated from the bottom of the screen to their correct position at the top of the screen (see Figure 10). A line was drawn across the middle of the screen. After the letters passed this visible line in the middle of the screen, participants could no longer select and click on letters they thought to be incorrect. Like in the LA condition, participants were forced to make quick decisions.

Participants assigned to the diagram assist (DA) condition were charged with the same string-edit task as the other conditions, but were provided with a diagram of the finite-state

grammar (see figure 10) for assistance. Instead of the letters floating to the top of the screen as in the PA condition, the letters appeared, one by one, in the correct order and position in the state diagram from left to right. Also like the other two conditions, quick decisions were required; after a letter appeared in the diagram participants could no longer click on the corresponding letter in the string.

c x T x x _ _ _ _

[c] [x] [T] [x] [⋈] [T] [T] [v] [v]

_ _ _ _ _ _ _ _
  c x T s s s
        s s s x s

[c] [⋈] [⋈] [s] [s] [s] [s] [v] [x] [s]

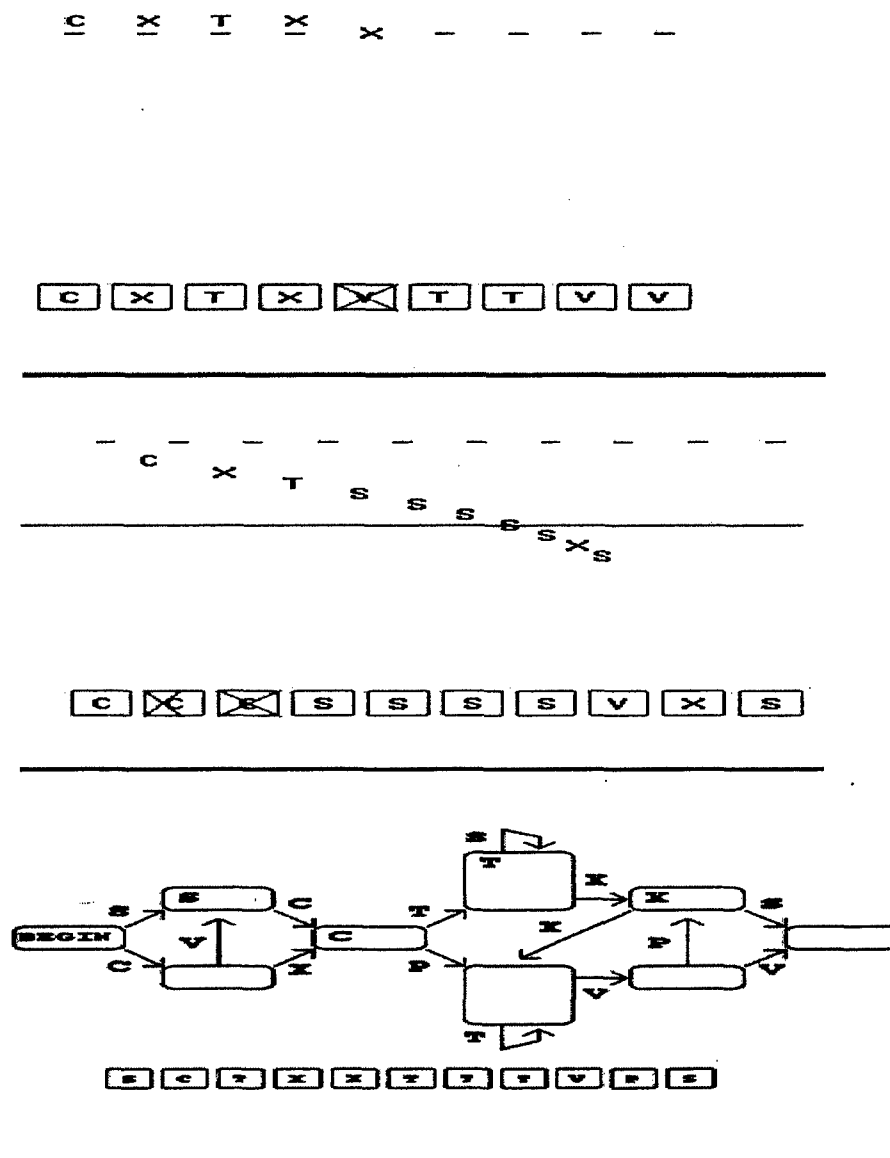[s] [c] [T] [x] [x] [T] [T] [T] [v] [P] [s]

Figure 10. *Screen shots from the middle of a trial in the three practice conditions, LA, PA, and DA, respectively*

56

Following training, a transfer test using the same cued-generate task used in the previous experiments was used to compare performance across conditions. Participants were required to generate exemplars based on two randomly selected cues.

Method

Participants. One-hundred and thirteen undergraduate psychology students taking a number of different courses at Louisiana State University participated in this experiment. All participants were volunteers and were compensated for their participation with extra credit.

Materials. This experiment used the same finite-state grammar from Mathews et al. (1989). 177 letter strings, or exemplars, ranging from 5 to 11 letters in length are generated by this grammar. A subset of 22 exemplars was randomly selected by the computer at the beginning of each training phase, for each participant. Each exemplar was seen approximately 4 times in the training phase.

Design. The design was a one-factor between subjects design with four levels: letter appearance, primed assist, diagram assist, and the no-training control. Subjects were randomly assigned to each of the four groups. Attrition among subjects caused the groups to be of unequal size; LA (27 participants), PA (28 participants), DA (24 participants), and control (34 participants).

Procedure. Participants were tested in groups up to 5. Each participant attended three sessions over the course of one week. Data from subjects who did not attend all three sessions were not included in the analysis. Sessions one and two began with a 20-min training phase requiring participants to perform their assigned training task. The training phase was followed by a 20 min cued-generate test phase. A retention test was given without a training phase on session three.

Testing Phase. During the cued-generate task, the computer displayed a set of dashes, corresponding to the number of letters in the target letter string. Two randomly selected letters from a no-yet-generated exemplar were displayed on two of the dashes. Participants filled in each blank dash with a letter and pressed enter. If the letter string generated by the participant did not match at least 70% of the letters in the closest not-yet-generated string, the participant was required to make another attempt. If any letters matched a not-yet-generated string, they were displayed on this new attempt, along with the two cues from the first attempt. The participant repeated this process until they had matched at least 70% of the letters. When the participant reached the 70% criterion, the letter string that they created was displayed along with the target string and the percent of letters matched. Each exemplar could only be generated once per session.

Participants in the test-only control were given the six letters randomly typed across the middle of a page. The control participants were given the same codeword cover story and instructions as the other groups.

Participants were instructed to work as quickly as possible while still being accurate. A monetary prize was offered to the participant in each condition who generated the most exemplars across all three sessions.

Results

    With the exception of the training task error data, only data from session three were analyzed as this was the only session in which a training phase did not immediately precede the test phase. The results for all five dependent measures are presented in Table 8. The results for each measure are discussed.

	Errors (session 2)	Achievement	Efficiency	Perfects	Speed
Control	---	0.1044 (.0182)	0.0388 (.0073)	0.0053 (.0006)	10.478 (.4641)
LA	2.0118 (.1782)	1.3389 (.1881)	0.3884 (.0438)	0.0413 (.0091)	7.3315 (.4553)
DA	1.1769 (.2405)	0.9479 (.2474)	0.2973 (.0677)	0.0936 (.0448)	7.618 (.6694)
PA	0.249 (.0210)	0.9268 (.1904)	0.2511 (.0454)	0.0147 (.0036)	9.07 (.5727)

Table 8. *Means and Standard Error (in parentheses) of Artificial Grammar Experiment 4 for Final Test*

    Training Errors. Training errors were measured in terms of the number of hits and false alarms per trial in the string-edit task of the training phase. A repeated-measures analysis of variance (ANOVA) was used to analyze these data. There was a significant effect of sessions, $F(1, 78) = 45.067$, $p < .001$. There was also a significant effect of group in session one, $F(2, 76) = 38.346$, $p < .001$. A Tukey Honestly Significantly Different (HSD) post hoc test of comparisons showed that the PA group ($M=.439$) made significantly fewer errors than the DA ($M=1.944$) and the LA ($M=2.461$).

    Session two showed similar results. Again, there was a significant effect of group, $F(2, 76) = 29.80$, $p < .001$. A Tukey HSD post hoc test of comparisons showed that the PA group ($M=.249$) made significantly fewer errors than the DA group ($M=1.18$), which made significantly few errors than the LA group ($M=2.01$).

    Achievement. We measured achievement as the number of acceptable strings (those matching at least 70% of the letters from the target exemplar) generated *on the first attempt* at each target exemplar per minute of the 20-min test phase. Note that this measure of achievement is different from the way achievement was measured in Experiments 1-3. We changed this measure because the control group simply pressed keys rapidly without knowing anything about the correct strings. With this measure such random key pressing will not result in high achievement. An attempt was recorded each time a participant filled in the blanks and pressed the enter key. A one-way analysis of variance (ANOVA) was used to analyzed the data. There was a significant effect of group $F(3, 109) = 10.609$, $p < .001$. A Tukey HSD post hoc test of

comparisons showed that the test-only control group (M=.014) showed significantly lower achievement than the LA (M=.1.34), DA (M=.948) and PA (M=.927) groups, which did not differ significantly.

Efficiency. Efficiency was measured as the proportion of first attempts on a target exemplar that generated acceptable strings. A one-way analysis of variance (ANOVA) was used to analyze the data. There was a significant effect of group, $F(3, 109) = 13.51$, $p < .001$. A Tukey HSD post hoc test of comparisons showed that that the LA (M=.388), DA (M=.297), and PA (M=.25) groups were significantly more efficient than the control (M=0.039).

Perfects. Perfects were a measure of the proportion of letter strings generated on the first attempt that matched 100% of the letters in the target exemplar. A one-way analysis of variance (ANOVA) was used to analyze the data and showed a significant effect of group, $F(3, 109) = 3.87$, $p < .05$. A Tukey HSD post hoc test of comparisons showed that that the DA (M=.094) produced more perfect strings than the PA (M=.0147) and the control (M=.0053) groups. The LA group (M=.041) did not differ significantly from any group.

Speed. Speed was a measure of the number of attempts made per minute. A one-way analysis of variance (ANOVA) was used to analyzed the data. There was a significant effect of group, $F(3, 109) = 7.796$, $p < .001$. A Tukey HSD post hoc test of comparisons showed that that the test-only control (M=10.48) responded at a significantly higher speed than the DA (M=7.62) and LA (M=.7.33) groups. The PA group (M=9.07 did not differ significantly from any group.

Discussion

During the training phase, each group performed the same string edit task, but with a different type of assistance. Participants in the LA condition just attempted to identify wrong letters in the strings. The AD condition did the same editing task, but was provided with a state diagram of the grammar for assistance. Finally, the PA condition was aided by the letters rising to the top of the screen to prime the correct choices of letters in the string to edit. The PA group far outperformed the other groups in the editing task, but did not transfer that superior performance to the sting generation task.

All three groups receiving training had higher achievement and were more efficient than the test-only control. The number of perfect strings generated on the first attempt did differ among the trained groups. The DA condition generated more perfects than the PA and control conditions. While the DA condition generated nominally more perfect strings than the LA condition, the difference was not significant. This is likely due to the large degree of within group variability. It is possible that this within group variability may be decreased with more training. This is explored in experiment 5.

Introduction

One typical criticism, going back to Miller (1968), of the using an artificial grammar paradigm to study generativity is that participants are exposed to the grammar for a very short period of time. In the current experiment, participants completed four 20-min training phases across two weeks, compared with two 20-min session in experiment 4. Also, participants saw four randomly generated 20-exemplar sets from the finite-state grammar, compared with two sets from experiment 4.

Additionally, participants took two types of tests in the current experiment. The tests were divided into two parts, a speed portion and an accuracy portion. During the speed test, participants were given 10-mins to make as many attempts as they wished. In the accuracy test, participants were allowed 60 attempts and were encouraged to contemplate their responses as there was no time limit. Both tests were administered on the same day (during sessions three and six), with the speed test followed by the accuracy test.

We expected that all participants would change their strategies in relation to the type of test given (i.e. respond slowing but accurately in the accuracy test and respond quickly in the speed test). Additionally, we expected that the DA group would perform better under the slow pace encouraged by the accuracy test. Conversely, it was thought that the PA group would perform better in the speed test, where a fast pace was encouraged.

Method

Participants. Eighty undergraduate psychology students taking a number of different courses at Louisiana State University participated in Experiment 4. All participants were volunteers and were compensated for their participation with extra credit

Materials. The same materials from Experiment 4 were used in the current experiment.

Design. The design was a one-factor between-within-subjects design with four levels between subjects: LA condition, PA condition, DA condition, and the no-training control condition. The two within-subject factors were test type: speed test and accuracy test. Subjects were randomly assigned to each of the four groups. Attrition among subjects caused the groups to be of unequal size; LA (20 participants), PA (21 participants), DA (20 participants), and control (19 participants).

Procedure. Participants were tested in groups up to 5. Each participant attended six sessions over the course of two week. Data from subjects who did not attend all six sessions were not included in the analysis. Sessions one, two, four, and five began with a 20-min training phase requiring participants to perform their assigned training task. The training phase was followed by a 20 min cued-generate test phase. The speed and accuracy tests were given without a training phase on sessions three and six.

The same code-word cover story used in Experiment 1 was used in the current research and participants were told that a monetary prize would be given to whomever made the fewest errors in the training phase and found the most code-words in the test phase.

Results

With the exception of the training task error data, only the data from six were analyzed, as that is where the treatment effect was strongest. The results for all five dependent measures are presented in Table 9.

Table 9.
Results from Final Test Performance in Artificial Grammar Experiment 5

	Errors (session 5)	Achievement	Efficiency	Perfects	Speed
Control speed test	---	1.04	0.227	0.0189	4.0579
Control accuracy test		1.299	0.2854	0.0245	3.8118
LA speed test	1.8077	2.07	0.4194	0.0654	4.46
LA accuracy test		2.16	0.4752	0.0791	3.9267
DA speed test	0.8263	3.14	0.5488	0.238	5.095
DA accuracy test		3.0	0.5845	0.2633	4.6358
PA speed test	0.2579	2.96	0.5332	0.051	5.12
PA accuracy test		3.244	0.5845	0.0762	5.177

Training Errors. A repeated-measures analysis of variance (ANOVA) was used to analyze these data. There was a significant effect of sessions, with errors decreasing from session one to session four, $F(3, 159) = 24.605$, $p < .001$. This factor did not interact with group, $F(6, 159) = 1.012$, ns. There was also a significant of group at each session. The pattern of results was similar across all four sessions, so we will present only the mean error score for each group across all four sessions. $F(2, 55) = 12.151$, $p<.001$. A Tukey HSD post hoc comparison showed that the PA group (M=.26) and DA group (M=.83) made significantly fewer errors than the LA group (M=1.8).

Achievement. A one-way analysis of variance (ANOVA) was used to analyze the speed test data. There was a significant effect of group $F(3, 74) = 5.106$, $p < .01$. A Tukey HSD post hoc test of comparisons showed that the DA (M= 3.14) and PA (M=2.94) groups had higher achievement than the test-only control (M=1.29). The LA group (M=2.16) did not differ significantly from any other group.

The accuracy test data were analyzed in the same manner, and showed a significant effect of group, $F(3, 74) = 4.199$, $p < .05$. A Tukey HSD post hoc test of comparisons showed the same pattern as in the speed test. The DA (M=3.0) and PA (M=3.24) groups had higher achievement than the test-only control (M=1.29). Again, the LA group (M=2.16) did not differ significantly from any other group.

A repeated measures ANOVA with test type as the repeated factor and group as the between factor showed no significant difference between achievement on the speed (M=2.34) and accuracy (M=2.45) tests, $F(1, 76) = 1.69$, ns. No significant interaction between group and

61

test type was found either, F(3, 76) = 1.856, ns.

Efficiency. A one-way analysis of variance (ANOVA) was used to analyze efficiency during the speed test. There was a significant effect of group F(3, 74) = 5.032, p < .01. A Tukey HSD post hoc test of comparisons showed that the DA (M= 0.55) and PA (M= 0.53) groups had higher achievement than the test-only control (M= 0.23). The LA group (M= 0.43) did not differ significantly from any other group.

Efficiency data from the accuracy were analyzed in the same manner as the speed test. Again, there was a significant effect of group, F(3, 74) = 3.967, p < .05. Tukey HSD post hoc comparison showed that the DA (M = 0.59) and PA (M = 0.57) groups had higher efficiency than the test-only control (M=0.29). The LA group (M= 0.48) did not differ significantly from any other group.

A repeated measures ANOVA with test type as the repeated factor and group as the between factor showed significant difference between achievement on the speed (M= 0.44) and accuracy (M=0.49) tests, F(1, 76) = 9.173, p < .01. No significant interaction between group and test type was found, F(3, 76) = 0.178, ns.

Perfects. There was a significant effect of group on the speed test data, F(2, 74) = 6.252, p<.01. Tukey HSD post hoc procedures showed that the DA group (M=0.24) had a significantly higher number of perfect entries on their first attempt than the PA (M=0.063), LA (M=0.056), and test-only control (M=0.019) groups. No other pairwise comparisons were significant.

There was also a significant effect of group in the accuracy test, F(3, 74) = 7.318, p < .001. Tukey HSD post hoc procedures revealed the same pattern as in the speed test. The DA group (M=0.26) had a significantly higher number of perfect entries on their first attempt than the PA (M=0.085), LA (M=0.08), and test-only control (M=0.025) groups. No other pairwise comparisons were significant.

A repeated measures ANOVA with test type as the repeated factor and group as the between factor showed significant difference between number of perfects on the speed (M= 0.096) and accuracy (M=0.116) tests, F(1, 76) = 8.339, p < .01. No significant interaction between group and test type was found, F(3, 76) = 0.480, ns.

Speed. A one-way ANOVA found a significant effect of group on the speed test, F(3, 74) = 3.137, p < .05. Tukey HSD post hoc procedures revealed that the test-only control group (M=10.93) made more attempts than the PA group (M=8.45). Pairwise comparisons involving the DA (M=8.86) and LA (M=9.22) groups were not significant.

There was also a significant effect of group on the accuracy test. F(3, 74) = 3.14, p < .05. Pairwise comparisons using Tukey post hoc procedures showed that the test-only control (M=9.59) made more attempts than the LA group (M=7.4). Pairwise comparisons involving the DA (M=7.9) and PA (M=8.13) groups were not significant.

A repeated measures ANOVA with test type as the repeated factor and group as the between factor showed significant difference between number of attempts made on the speed (M=9.3) and accuracy (M=8.2) tests, F(1, 76) = 33.391, p < .001. No significant interaction between group and test type was found, F(3, 76) = 2.526, ns.

Discussion

The results of Experiment 5 follow those from Experiment 4 with one important

exception. The DA group, which combined experience with exemplars and knowledge of the grammar's structure had a greater number of perfect responses than all other conditions. This shows that with a lengthy training phase, model-based processing can result in performance that is as fast as memory-based processing while at the same time being more accurate.

No differences were found between on any of the dependent measures between the speed and accuracy tests. It is possible that with the large amount of training our participants received, they were so accustomed to using one strategy, that they were unable to switch strategies when the demands of the task changed. Or, they may have felt that their strategy for one test was also appropriate for the other and there was no need to switch.

General Discussion

Now it is time to have a general discussion of the studies in both process control and artificial grammars domains, including results from both human experiments and computational simulations. Below we will highlight a few points that we consider particularly important or prominent from our studies.

In the process control task, learners appear to acquire correct responses more from implicit induction rather than explicit rule generation. In fact, our college level participants were particularly bad at figuring out the relatively simple equation that determined reactor temperature, even when they were assisted by giving hints. Yet a simple cue that includes three good examples, such as, "If current temperature is 10,000 then use 800 pellets", was effective in enhancing learning. Any form of concurrent reflection was found to have a negative impact on learning in this paradigm. However, post task reflection was somewhat beneficial early in learning.

Five experiments contrasted grammar learning following various combinations of purely implicit training (EP), purely explicit training (GR) and integrated training (ED). Implicit training consisted of copying exemplars into a response sheet that required attention to the serial order of letters in each string. Purely explicit training consisted of memorizing a transition diagram of the grammar. The integrated training consisted of copying exemplars into the transition diagram. The cued-generate test was used to test the ability of participants to generate a wide range of grammatical strings under conditions where perfect performance was not required (70% match to a valid string was acceptable).

The overall pattern of results can be summarized very simply: Implicit training led to fast but relatively inaccurate generation of strings and explicit training led to very slow but relatively accurate string generation.

The notion that implicit learning is very inflexible was not supported. Groups that only received implicit training (the EP task) performed very well on the cued-generate test in terms of total number of strings generated (the achievement measure). In fact, in Experiment 3, the implicit group successfully generated nearly twice as many strings as the purely explicit trained group (the GR group).

Experiments 2 and 3 on artificial grammars also provide interesting findings concerning attempts to mix the two types of training. Two types of mixing were employed. In some cases purely implicit training and purely explicit training were switched across two sessions (Experiment 3). In other groups implicit and explicit training were integrated into one type of training exposing participants to exemplars and mapping their structure onto a diagram of the grammar (the ED task). Surprisingly, the integrated training did not lead to greater achievement than purely implicit training. Also, in Experiment 2, both groups that received integrated training (ED) in one of the two sessions and implicit training in the other (EP,ED or ED,EP), ended up showing the relatively fast but inaccurate performance associated with purely implicit learning. This pattern suggests that participants preferred the implicit mode when exposed to both implicit and integrated training.

The results were a bit different in Experiment 3 where purely implicit (EP) training was mixed across sessions with purely explicit (GR) training. Regardless of training order, both of

64

these mixed groups showed the increased accuracy and slower speed associated with explicit learning on Session 2. However, the group that got explicit training first (GR,EP) did not reach the achievement level associated with purely implicit training or implicit followed by explicit training (EP,GR). In some sense this group got the worst of both types of training - they ended up being relatively slow and inaccurate. It is a bit alarming that this pattern of training (explicit followed by implicit) might best characterize training outside the laboratory. This would be the pattern associated with formal schooling (explicit training) followed by experience with many cases (implicit training) when one gets a job.

Another question raised by this research concerns the tendency to prefer the implicit mode when exposed to a mixture of integrated and implicit training. Perhaps a similar phenomenon would occur outside the laboratory when people are explicitly trained in school and then practice on their job. That is, there might be a tendency to move toward the implicit mode as one gains experience and this shift might lead to decreased accuracy of judgment. One recent research of radiologists (Beam, Conant, & Sickles, 2003) supports such a decrease in accuracy in performance associated with practice following completion of formal education. This research found a small but significant drop in cancer detection for each year beyond a doctor's residency training. We are currently planning experiments to explore this possibility.

Experiment 5 found that the best of both worlds could be achieved by using an animated diagram of the grammar to prime learning during practice. In this case the group that processed exemplars while simultaneously seeing the string diagramed in the animation achieved both high speed and accuracy. The use of explicit structure to enhance rather than compete with implicit learning appears to be a promising path for more research.

The practical messages of this research for training are straightforward: If only accuracy matters use explicit training. If only speed counts, use implicit training. If both speed and accuracy are important the mixed training may be best. The best results were obtained in using an animated diagram of the grammar appearing while learners were concentrating on finding and correcting errors quickly in whole grammar strings. The difference in this form of training from the integrated form of training used in the earlier experiments (which produced good but not best levels of performance) appears to be related to the emphasis on speed in training and having learners' attention focused on synthesizing whole strings (the implicit mode) rather than analyzing strings into the diagram (the explicit mode). We believe that this emphasis on quickly and implicitly processing whole strings and using the explicit structure (diagram) to help understand how strings are made facilitates memory-based implicit processing that is essential for implicit learning (Domangue, Mathews, Sun, Roussel, & Guidry, 2004).

The above points are being verified through computational simulations using the CLARION cognitive architecture. Moreover, the CLARION simulation of process control has led us to formulate and test those hypotheses concerning process control learning in the first place. At the same time, discrepancies between theoretical models and experimental data have led to new designs of further human experiments. It is particularly important that simulations in various domains, ranging from process control and artificial grammars to Tower of Hanoi, all indicated the importance of the implicit/explicit interaction in enhancing skill acquisition and training. Specifically, the findings that explicit processing should complement but not compete with implicit processing have been confirmed by simulations in process control, artificial grammar, minefield navigation, Tower of Hanoi, and a variety of other domains (Sun 2002, Sun

et al 2001, Sun and Zhang 2003, 2004). It appears that there is a useful lesson that can be drawn from all the studies above, including both human experiments and computational modeling and simulation.

Summary and Conclusions

Let us summarize the work in both process control and artificial grammars domains, including summarizing the results from both human experiments and computational simulations, and draw a few general conclusions.

The current work advances basic research in the areas of learning and cognition. One product of this effort is a conceptual framework, which addresses the ways these two types of knowledge interact to produce expertise (e.g., in tasks that require both speed and accuracy), which is an open, but important, issue. This framework (the CLARION cognitive architecture) suggests that human performance may be controlled by either a subconceptual knowledge base (the implicit mode) or application of a symbolic conceptual mental model (the explicit mode). Implicit control is fast but prone to error, particularly in early levels of skill acquisition. Explicit control is more accurate but slow to apply, and prone to loss by forgetting over a retention interval. We have found that reflection about how one is performing the task can be beneficial following short periods of practice. However, it is often even more effective when learners are provided hints that direct their reflection in productive directions. These are important findings that advance our understanding of the interaction of the two types of knowledge.

A computational cognitive architecture, CLARION, significantly different from other existing cognitive architectures, is developed in this work to simulate and capture a range of quantitative data that are related to the interaction, based on the above ideas. This will help us to capture and explain (and eventually to predict) training and learning processes. We carry out simulation experiments in the domains of process control tasks, artificial grammar learning tasks, as well as many other tasks (Sun 2002), and generate new insight and interpretations that can further explicate the interaction between implicit and explicit processes. These outcomes (data, models, and theories) provide a more detailed, clearer and more comprehensive perspective on skill learning. Our models and theories will be useful in better understanding human skill learning, as well as in helping to improve learning processes. Our models and theories may also be useful in understanding individual differences in skill learning (based on the implicit/explicit interaction). Since the CLARION cognitive architecture and simulations based on it have been published in many journal papers and books (see, e.g., Sun 2002, Sun et al 2001), we did not describe most of them, except highlighting two most relevant simulations earlier.

The results of our experiments support our theory/model of the interactions of implicit and explicit learning processes during skill acquisition. Strictly implicit training is effective for fast responding, but is prone to error. Strictly explicit training results in slow but accurate responding. A balance of both worlds (fast and accurate responding) can be obtained by using structural models in training that emphasize fast but accurate responding. Under these training conditions, learners acquire the ability to rely on implicit knowledge for generating an initial sketch of a solution and using explicit knowledge to fill in gaps or check possible errors.

Our research also demonstrates that implicitly acquired knowledge can be much more flexible than existing research suspected. It was believed that implicitly acquired knowledge would not generalize beyond experienced cases. However, we found that people could acquire knowledge from artificial grammar cases that could be recombined to generate a range of valid strings not yet experienced. This form of learning would be especially valuable if combined with

external help that could correct minor errors, such as the computer did for our participants during the artificial grammar generation task (e.g., whenever a response reached an acceptable level it was corrected by the computer).

In the process control task, learners appear to acquire correct responses more from implicit induction rather than explicit rule generation. In fact, our college level participants were particularly bad at figuring out the relatively simple equation that determined reactor temperature. Yet a simple cue that includes three good examples, such as, "If current temperature is 10,000 then use 800 pellets", was effective in enhancing learning. Perhaps these hints showed learners how to look at the task in terms of finding good cases. Reflective thinking in between practice sessions did enhance performance early in learning. Therefore, the type of training recommended for this type of complex process control task consists of short periods of fast, intense practice followed by short intervals of reflection. Also providing learners of a few examples of good responses to specific situations can be very effective.

These results should be further developed, because they may have significant implications for Army training and for other applied areas of the Army. The knowledge gained from the basic research would apply when developing training programs, with a better understanding of the cognitive processes involved in skill acquisition, both implicit and explicit, and when addressing how to increase training effectiveness. In particular, this basic research program addresses an important issue when developing training programs, how implicit and explicit processes interact and impact skill learning and performance. Much more work is needed in this area. Similarly, focus of research in decision making has been on cognitive skills training methods that facilitate rapid, accurate decision-making. Although these are different foci, this basic research could inform the applied research on important considerations in decision-making.

Some of the above hypotheses have been verified through computational simulation using CLARION. In particular, the CLARION simulation of learning process control has led us to formulate and test those hypotheses concerning process control learning. At the same time, discrepancies between theoretical models and experimental data have led to new designs of further human experiments. It appears that CLARION has the potential to be a comprehensive theory of a range of psychological tasks/domains (Sun 2002). Future research should be conducted to further develop and validate this approach.

REFERENCES

Ahlum-Heath, M. E., & DiVesta, F. J. (1986). The effects of conscious controlled verbalization of a cognitive strategy on transfer in problem solving. *Memory and Cognition, 14,* 281-285.

Anderson, J. R. (1993). *The rules of the mind.* Lawrence Erlbaum Associates. Hillsdale, NJ.

Beam, C. A., Conant, E. F., & Sickles, E. A. (2003). Association of volume-independent factors with accuracy in screening mammogram interpretaton. *Journal of the National Cancer Institute, 95,* 282-290.

Berry, D. C. (1983). Metacognitive experience and transfer of logical reasoning. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 35,* 39-49.

Berry, D. C., & Broadbent, D. E. (1984). On the relationship between task performance and verbalizable knowledge. *Quarterly Journal of Experimental Psychology, 36,* 209-231.

Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and implicit-explicit distinction. *British Journal of Psychology, 79,* 251-272.

Brooks, L. R. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169-211). Hillsdale, NJ: Erlbaum.

Chan, C. (1992). Implicit cognitive processes: Theoretical issues and applications in Computer systems design. D. Phil: thesis, University of Oxford.

Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self explanations: How students study and use examples in learning to solve problems. *Cognitive Science, 13,* 145-182.

Chi, M. T. H., DeLeeuw, N., Chiu, M., & LaVarcher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science, 18,* 439-477.

Clark, A & Karmiloff-Smith, A. (1993). The cognizer's innards. *Mind and Language, 8,* 487-519.

Cleeremans, A. (1994). Attention and awareness in sequence learning. *Proeedings In Cognitive Science Society Annual Conference,* 330-335.

Dienes, Z., & Altmann, G. (1997). Transfer of implicit knowledge across domains: How implicit and how abstract? In D. Berry (Ed.), *How implicit is implicit learning?* (pp.107-123). Oxford, England: Oxford University Press.

Dienes, Z., & Berry, D. C. (1997). Implicit synthesis. *Psychonomic Bulletin & Review, 4*, 68-72.

Dienes, Z., Broadbent, D. E., & Berry, D. C. (1991). Implicit and explicit knowledge bases in artificial grammar learning, *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*, 875-887.

Dienes, Z., & Fahey, R. (1995). Role of specific instances in controlling a dynamic system. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 848-862.

Domangue, T. J., Mathews, R. C., Sun, R., Roussel, L. G., and Guidry, C. E. (2004). Effects of model-based and memory-based processing on speed and accuracy of grammar string generation. *Journal of Experimental Psychology: Learning Memory, and Cognition: 30*, 1002-1011.

Dulany, D. E., Carlson, R. A., & Dewey, G. I. (1984). A case of syntactical learning and judgment: How concrete and how abstract? *Journal of Experimental Psychology: General, 113*, 541-555.

Gibson, F., Fichman, M., & Plaut, D. (1997). Learning in dynamic decision tasks. *Organizational Behavior and Human Decision Processes, 71*, 1-35.

Howard, J. H., & Ballas, J. A. (1980). Syntactic and semantic factors in the classification of nonspeech transient patterns. *Perception and Psychophysics, 28*, 431-439.

Hunt, E., & Lansman, M. (1986). Unified model of attention and problem solving. *Psychological Review, 93*, 446-461.

Knowlton, B. J., & Squire, L. R. (1996). Artificial grammar learning depends on implicit acquisition of both abstract and exemplar-specific information. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22*, 169-181.

Lebiere, C, Wallach, D., Taatgen, N. (1998). Implicit and explicit learning in ACT-R. *Proeddings of ECCM'98*, 183-189. Nottingham University Press.

Lane, S. M. & Schooler, J. S. (in press). Skimming the surface: The verbal overshadowing of analogy retrieval. *Psychological Science*.

Lee, Y. (1995). Effects of learning contexts on implicit and explicit learning. *Memory and Cognition, 23*, 723-734.

Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory & Cognition, 13*, 523-530.

Manza, L., & Reber, A. S. (1997). Representing artificial grammar: Transfer across stimulus forms and modalities. In D. Berry (Ed.), *How implicit is implicit learning?* (pp. 73-106). Oxford, England: Oxford University Press.

Mathews, R. C. (1991). The forgetting algorithm: How fragmentary knowledge of exemplars can abstract knowledge. *Journal of Experimental Psychology: General, 120*, 117-119.

Mathews, R. C. (1997). Is research painting a biased picture of implicit learning? The dangers of methodological purity in scientific debate. *Psychonomic Bulletin & Review, 4*, 38-42.

Mathews, R. C., Buss, R. R., Stanley, W. B., Blanchard-Fields, F., Cho, J. R., & Druhan, B. (1989). The role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 1083-1100.

Mathews, R. C., Roussel, L. G., Cochran, B. P., Cook, A. E., & Dunaway, D. L. (2000). The Role of implicit learning in the acquisition of generative knowledge. *Journal of Cognitive Systems Research, 1*, 161-174.

Mathews, R. C., Lane, I. M., Roussel, L. G., Nagy, M. S., Haptonstahl, D. E., & Brock, D. B. (1996). Using conscious reflection, group processes and AI to facilitate development of expertise. *Journal of Experimental and Theoretical Artificial Intelligence, 8*, 259-276.

Meissner, C. A., Brigham, J.C., & Kelley, C. M. (2001). The influence of retrieval processes in verbal overshadowing. *Memory & Cognition, 29*, 176-186.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limit on our capacity for processing information. *Psychological Review, 63*, 81-97.

Miller, G. A. (1969). *The psychology of communication: seven essays*. Basic Books, New York.

Moscovitch, M., & Umilta, C. (1991). Conscious and unconscious aspects of memory. *Perspectives on cognitive neuroscience*. Oxford University Press, New York.

Norman, D. A. (1993). *Things that make us smart*. Reading, MA: Addison-Wesley.

Nosofsky, R, Palmeri, T., & McKinley, S. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101*, 53-79.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General, 119,* 264-275.

Pollock, B. J., & Lee, T. D. (1997). Dissociated contextual interference effects in children and adults. *Perceptual and Motor Skills, 84,* 851-858.

Proctor, Robert W.; Dutta, Addie; (1995). *Skill acquisition and human performance.* Thousand Oaks, CA, US: Sage Publications, Inc, 1995. xvi, 442

Rabinbowitz, M., & Goldberg, N. (1995). Evaluating the structure-process hypothesis. In Weinert, F. & Schneider, W. (Eds.) *Memory Performance and Competencies.* Lawrence Erlbaum Associates. Hillsdale, NJ.

Reber, A. S. (1976). Implicit learning of synthetic languages: The role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 88-94.

Reber, A. S. (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General, 118,* 219-235.

Reber, A. S., Kassin, S., Lewis, S., & Canton, G. (1980). On the relationship between implicit and explicit modes in the learning of a complex rule structure. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 492-502.

Rosenbloom, T., Laird, J., & Newell, A. (1993). *The SOAR papers: Research on integrated intelligence.* MIT Press. Cambridge, MA.

Roussel, L. G. (1999). *Facilitating knowledge Integration and Flexibility: The effects of reflection and exposure to alternative models.* Unpublished doctoral dissertation. Louisiana State University.

Rumelhart, D., McClelland, J., & PDP research group (1986). *Parallel distributed processing.* MIT Press. Cambridge, MA.

Schacter, Daniel L. (1990). Introduction to "Implicit memory: Multiple perspectives." *Bulletin of the Psychonomic Society, 28(4),* Jul 1990. pp. 338-340.

Schooler, J.W., & Engstler-Schooler, T.Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology, 17,* 36–71.

Schooler, J. W., Ohlsson, S., & Brooks, K. (1993). Thoughts beyond words: When language overshadows insight. *Journal of Experimental Psychology: General, 122,* 166-183.

Seger, C. (1994). Implicit learning. *Psychological Bulletin, 115,* 163-196.

Servan-Schreiber, E., & Anderson, J. R. (1990). Learning artificial grammars with competitive chunking. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 592-608.

Shanks, David R. (1993).; Associative versus contingency accounts of category learning: Journal of Experimental Psychology: Learning, Memory, & Cognition, Vol 19(6), Nov 1993. pp. 1411-1423.

Shewokis, P. A., Del Ray, P., & Simpson, K. J. (1998). A test of retroactive inhibition as an explanation of contextual interference. *Research Quarterly for Exercise and Sport*, 69, 70-74.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences, 11*, 1-74.

Stadler, Michael A.; Warren, Justin L.; Lesch, Shana L.(2000). Is there cross-format transfer in implicit invariance learning? *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 53A(1), pp. 235-245.

Stanley, W. B., Mathews, R. C., Buss, R. R., & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task. *Quarterly Journal of Experimental Psychology, 41*, 553-577.

Sun R. (1994). *Integrating rules and connectionism for robust commonsense reasoning*. John Wiley and Sons, New York, NY, 1994.

Sun, R. (1995). Robust reasoning: Integrating ruled based and similarity based reasoning. *Artificial Intelligence, 75*, 241-296.

Sun, R. (1997). Learning, action, and consciousness. *Neural Networks*, 10, 1317-1331.

Sun, R. (2002). *Duality of the mind._* Lawrence Erlbaum Associates, Mahwah, NJ.

Sun, R. (2003). *A Tutorial on CLARION 5.0*. http://www.cogsci.rpi.edu/~rsun/sun.tutorial.pdf

Sun, R., Merrill, E., & Peterson, T. (1998). A bottom-up model of skill learning. *Proc. of 20th Cognitive Science Society Conference*, 1037-1042. Lawrence Erlbaum Associates, Mahwah, NJ.

Sun, R., Merrill, E., & Peterson, T. (2001). From implicit skill to explicit knowledge: A bottom-up model of skill learning. *Cognitive Science, 25*, 203-244.

Sun, R., & Peterson, T. (1998). Autonomous learning of sequential tasks. *IEEE Transactions on Neural Networks, 9,* 1217-1234.

Sun, R., & Zhang, X. (2003). Accessibility versus action centeredness in the representation of cognitive skills. *Proc. of the 5th International Conference on Cognitive Modeling,* 195-200. Bamberg, Germany.

Sun, R, & Zhang, X. (2004). Top-down versus bottom-up learning in cognitive skill acquisition. *Cognitive Systems Research, 5(1),* 63-89, March 2004.

Vokey, J. R., & Brooks, L. R. (1992). Salience of item knowledge in learning artificial grammar. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 328-344.

Watkins, C. (1989). *Learning with Delayed Rewards.* PhD Thesis, Cambridge University, Cambridge, UK.

Whittlesea, B. W., & Dorken, M. D. (1993). Incidentally, things in general are particularly determined: An episodic-processing account of implicit learning. *Journal of Experimental Psychology: General, 122,* 227-248.

Whittlesea, B. W., & Wright, R. L. (1997). Acting adaptively without knowing the consequences. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 23,* 181-200.

Willingham, D., Nissen, M., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 1047-1060.

Wilson, T.D., & Schooler, J.W. (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology, 60,* 181–192.